

STABILITY OF JOIN THE SHORTEST QUEUE NETWORKS

BY MAURY BRAMSON¹

University of Minnesota

Join the shortest queue (JSQ) refers to networks whose incoming jobs are assigned to the shortest queue from among a randomly chosen subset of the queues in the system. After completion of service at the queue, a job leaves the network. We show that, for all nonidling service disciplines and for general interarrival and service time distributions, such networks are stable when they are subcritical. We then obtain uniform bounds on the tails of the marginal distributions of the equilibria for families of such networks; these bounds are employed to show relative compactness of the marginal distributions. We also present a family of subcritical JSQ networks whose workloads in equilibrium are much larger than for the corresponding networks where each incoming job is assigned randomly to a queue. Part of this work generalizes results in [*Queueing Syst.* **29** (1998) 55–73], which applied fluid limits to study networks with the FIFO discipline. Here, we apply an appropriate Lyapunov function.

1. Introduction. Join the shortest queue (JSQ) refers to networks whose incoming “jobs” (or “customers”) are assigned to the shortest queue from among a randomly chosen subset of queues in the system. Here, shortest queue means the queue with the fewest jobs. (One could also consider the queue with the least remaining work.) Jobs are assumed to arrive in the network through one or more random streams of jobs. An arriving job is presented with a random subset B of the queues, with probability depending only on the stream, and chooses the queue in B with the fewest jobs; when two or more queues have the fewest jobs, one of these queues is chosen according to some rule. Jobs at queues are served according to some discipline, such as first-in, first-out (FIFO), last-in, first-out (LIFO), or processor sharing (PS) and, upon completion of service, leave the network.

A family of such networks is given by the *mean field* rule where $|B| = D$, $D \geq 1$, is fixed, and B is chosen uniformly from among the $\binom{N}{D}$ such sets, where N is the total number of queues. An alternative rule is given by choosing D queues with replacement from among the N queues, as in Vvedenskaya, Dobrushin and Karpelevich [21]. A natural setting for both rules is where all jobs arrive at a single Poisson stream with rate αN , with all jobs having the same service distribution $F(\cdot)$, with mean m , and are served according to the same discipline at each

Received April 2009; revised March 2010.

¹Supported in part by NSF Grants DMS-02-26245 and CCF-07-29537.

MSC2010 subject classifications. 60K25, 68M20, 90B15.

Key words and phrases. Join the shortest queue, stability.

queue. Such networks are *subcritical* when $\alpha m < 1$, that is, the long term rate at which jobs arrive in the network is strictly less than the total rate at which jobs are served when no queues are empty. When $\alpha m < 1$ and $F(\cdot)$ is exponentially distributed, it is elementary to show that the network is *stable*, that is, the underlying Markov process is positive Harris recurrent. It will therefore have an equilibrium distribution. In this setting, the choice of service discipline (when not relying on the residual service times) does not affect the distribution of the number of jobs at each queue, and the state space can be chosen so that it depends only on the number of jobs at each queue. When $D = 1$, we will say that each incoming job is *assigned randomly* to a queue; when, in addition, $F(\cdot)$ is exponentially distributed, the network will consist of N independent $M/M/1$ queues.

The asymptotic behavior of the equilibria of such networks, with D fixed as $N \rightarrow \infty$, has been studied since the mid 1990s. Let $E^{(N)}(\cdot)$ denote the equilibrium distribution function at a single queue for the network of N queues. In [21], it was shown that, for $\rho = \alpha m < 1$,

$$(1.1) \quad \lim_{N \rightarrow \infty} \bar{E}^{(N)}(\ell) = \rho^{1+D+\dots+D^\ell} = \rho^{(D^{\ell+1}-1)/(D-1)} \quad \text{for } \ell \in \mathbb{Z}_+,$$

where $\bar{E}^{(N)}(\cdot) = 1 - E^{(N)}(\cdot)$, and $D > 1$ is required for the second equality. Hence, the tail of $\lim_{N \rightarrow \infty} E^{(N)}(\cdot)$ decreases doubly exponentially fast when $D > 1$; when $D = 1$, the exponential tail is that of the corresponding $M/M/1$ queue. This rapid decrease in the tail has different applications, such as in the design of complex networking systems where memory is at a premium. See Azar et al. [1], Luczak and McDiarmid [12, 13], Martin and Suhov [14], Mitzenmacher [16], Suhov and Vvedenskaya [19], Vocking [20], and Vvedenskaya and Suhov [22] for related work on JSQ networks and ball-bin models in both theoretical and applied contexts.

The study of networks with given N has been more restricted. Foley and McDonald [10] studied the equilibria for small values of N .

Little work has been done on networks with nonexponential service distributions. In this setting, the stability of subcritical networks is no longer obvious. In particular, jobs might be assigned to short queues where the remaining work (or *workload*) is high, which can cause service inactivity after queues with many jobs, but low remaining work, empty. If the system can be “tricked” too often in this manner, it is conceivable that it is unstable while nevertheless being subcritical. For general service distributions, the evolution of the system will be influenced by the service discipline, which complicates analysis.

For JSQ networks with general service times, Foss and Chernova [11] is the main work that analyzes stability. Under the FIFO discipline, stability for a broad family of subcritical networks is demonstrated in [11], including those with the JSQ rule, for general service distributions and for arrivals given by a single renewal stream. Fluid limits are employed as the main tool. In this more general

framework, the appropriate definition for subcritical is no longer transparent. It will be discussed in the next subsection.

For our results, we adopt the same basic framework as in [11] for JSQ networks, but instead consider general service disciplines. We also allow multiple arrival streams. For general service disciplines, the number of partially served jobs may be large and known fluid limit techniques cannot be applied. Instead, we employ an appropriate Lyapunov function.

In this paper, we first show stability of subcritical networks for all nonidling disciplines. We then obtain uniform bounds on the tails of the marginal distributions of the equilibria for families of such networks; these bounds are employed to show relative compactness of the marginal distributions. Both the uniform bounds and relative compactness will be important tools for investigating, in the mean field setting, the limiting behavior, as $N \rightarrow \infty$, of the equilibria distributions $E^{(N)}(\cdot)$ at single queues for service disciplines such as FIFO, LIFO and PS (see Bramson, Lu and Prabhakar [6, 7]). Appropriate analogs of (1.1), for large values of ℓ , will hold under certain restrictions. We lastly present a family of subcritical mean field networks, with $N = D = 2$, where the service discipline is chosen so that the corresponding equilibria have much larger workload than do the corresponding $M/G/1$ queues. This shows that the JSQ rule does not always provide efficient service of jobs.

Main results. We present our main results here, Theorems 1.1–1.3, and discuss their ramifications. In the next subsection, we will give a general outline of the paper.

In order to avoid technical details, we postpone until Section 2 details regarding the construction of the state space S and of the Markov process $X(t)$, $t \geq 0$, underlying a JSQ network. We require at this point only limited specifics about the construction, namely that a state $x \in S$ is specified by descriptors that include the number of jobs z_n at each queue n , $n = 1, \dots, N$; the residual interarrival times u_k , $k = 1, \dots, K$, at each of the arrival streams, which are given by independent renewal processes; the residual service times $v_{n,i}$, $n = 1, \dots, N$ and $i = 1, \dots, z_n$, for each of the jobs currently in the network; and the ages $o_{n,i}$ for each of these jobs. (In the construction of S in Section 2, normalized versions of u_k , $v_{n,i}$ and $o_{n,i}$ are used.) When the arrival streams are Poisson, or when the service time distributions are the same and exponentially distributed, the corresponding descriptor may be dropped. The underlying Markov process $X(\cdot)$ takes values in S and is strong Markov.

In order to state our results, we need to specify the notion of subcriticality that was mentioned in the context of [11]. This requires the introduction of various terminology. We denote by $G_k(\cdot)$, $k = 1, \dots, K$, the distribution function for the interarrival time of jobs at the k th renewal stream, and by α_k the reciprocal of its mean, with $\alpha_k > 0$ being assumed. We denote by $p_{k,B}$ the probability that a job from arrival stream k chooses the shortest queue from the set B , $B \subseteq B_N$, where

$B_N = \{1, \dots, N\}$. We refer to B as the *selection set* and the rule corresponding to a given choice of $p_{k,B}$, $k = 1, \dots, K$, $B \subseteq B_N$, as the *selection rule*. For $n \in B$, such a job is a *potential arrival* at n .

We denote by $F_j(\cdot)$, with $j = (k, B, n)$ for $k = 1, \dots, K$, $B \subseteq B_N$ and $n \in B$, the distribution function for the service time (i.e., service requirement) of jobs from the k th renewal stream and selection set B that are served at queue n ; by m_j , the mean of $F_j(\cdot)$; and, by $\mu_j = 1/m_j$, the corresponding service rate. As in [11], we will require that either (a) $F_j(\cdot)$ depend only on n or (b) $F_j(\cdot)$ depend only on k and B , in which case we may write either $F_n(\cdot)$ or $F_{k,B}(\cdot)$ (and m_n or $m_{k,B}$, respectively, μ_n or $\mu_{k,B}$) when the context is clear. (When neither (a) nor (b) holds, analysis is more complicated and, as explained in Section 6 of [11], stability likely does not follow from subcriticality.) In [6, 7, 21], $F_j(\cdot) = F(\cdot)$ does not depend on k , B or N and there is a single renewal stream; in this setting, one can employ the notation α , m and μ . As in [11], we refer to networks satisfying (a) as *class independent* and (b) as *station independent*.

For the class independent case, we define the *traffic intensity*

$$(1.2) \quad \rho_1 = \max_{B \subseteq B_N} \left\{ \left(\sum_{n \in B} \mu_n \right)^{-1} \sum_k \sum_{A \subseteq B} \alpha_k p_{k,A} \right\}$$

and, for the station independent case,

$$(1.3) \quad \rho_2 = \max_{B \subseteq B_N} \left\{ |B|^{-1} \sum_k \sum_{A \subseteq B} \alpha_k p_{k,A} m_{k,A} \right\}.$$

When $\rho_1 < 1$, respectively, $\rho_2 < 1$, we say the network is *subcritical*. As was observed in [11], it is not difficult to check that when $\rho_i > 1$ in either (1.2) or (1.3), the corresponding network will be unstable. This behavior does not depend on the service discipline.

When the network is both class and station independent, (1.2) and (1.3) reduce to

$$(1.4) \quad \rho \stackrel{\text{def}}{=} \rho_1 = \rho_2 = \max_{B \subseteq B_N} \left\{ \frac{m}{|B|} \sum_k \sum_{A \subseteq B} \alpha_k p_{k,A} \right\}.$$

Let H be a subgroup of the permutation group on B_N on which all queues communicate [i.e., for given $n_1, n_2 \in B_N$, $\pi(n_1) = n_2$ for some $\pi \in H$], and assume the symmetry condition

$$(1.5) \quad \sum_k \sum_{A \subseteq B_\pi} \alpha_k p_{k,A} = \sum_k \sum_{A \subseteq B} \alpha_k p_{k,A} \quad \text{for all } B \subseteq B_N \text{ and } \pi \in H$$

is satisfied, where $B_\pi = \{n : n = \pi(n') \text{ for some } n' \in B\}$. (This holds, in particular, in the mean field setting.) Also, assume the network is both class and station independent. Then it is not difficult to check that (1.4) reduces to

$$(1.6) \quad \rho = \left(\sum_k \alpha_k \right) m / N.$$

(Note that, for given A , $A \subseteq B_\pi$ for at most $|H||B|/N$ permutations $\pi \in H$, with equality holding when A is a singleton.)

In addition to subcriticality, we will require the following condition on JSQ networks for Theorems 1.1–1.3: for some $\Gamma \in \mathbb{Z}_{+,0}$ and $h(\cdot)$, with $h(t) > 0$ for all t ,

$$(1.7) \quad P_x(\text{at most } \Gamma \text{ potential arrivals at } n \text{ occur over } (0, m^{\max}t]) \geq h(t)$$

for all x, t and n , where $m^{\max} \stackrel{\text{def}}{=} \max_j m_j$. Note that this condition depends only on the distributions $G_k(\cdot)$ and probabilities $p_{k,A}$. It is met in most cases, for instance, if (a) for each k and n , $\sum_{A \ni n} p_{k,A} < 1$, in which case one can set $\Gamma = 0$, or (b) for each k and y , $G_k(y) < 1$, where $\Gamma = K + 1$ suffices. Condition (a) always holds in the mean field setting when $D < N$; Condition (b) is equivalent to (1.11), which is required for Corollary 1.1.

In Theorems 1.1–1.3, we will employ the nonnegative function $\|x\|$, or *norm*, for $x \in S$. It is defined in terms of the norms $\|x\|_L$, $\|x\|_R$ and $\|x\|_A$ by

$$(1.8) \quad \|x\| = \|x\|_L + \|x\|_R + \|x\|_A.$$

We define these components of $\|x\|$ in Section 4. Without going into details here, we note that $\|x\|_L$ depends on the number of jobs and a truncation of the residual service time of each job; $\|x\|_R$ depends on just the residual service time of each job and will be employed for residual service times greater than the previous truncation; and $\|x\|_A$ measures the residual interarrival times with appropriate weighting. For given $M > 0$, we denote by $\tau_M(1)$ the stopping time

$$(1.9) \quad \tau_M(1) = \inf\{t \geq 1 : \|X(t)\| \leq M\}.$$

We now state Theorem 1.1. Here and elsewhere in the paper, we implicitly assume the discipline is nonidling. In Section 2, we will also specify mild conditions on how the service effort devoted to individual jobs is allowed to change over time.

THEOREM 1.1. *For each subcritical JSQ network satisfying (1.7), there exist M and C_1 so that*

$$(1.10) \quad E_x[\tau_M(1)] \leq C_1(\|x\| \vee 1) \quad \text{for all } x,$$

where $\|x\|$ is the norm given in (1.8).

For certain service disciplines, such as FIFO and PS, the condition (1.7) can be avoided; this is noted after the proof of Proposition 5.1. As mentioned above, (1.7) will in most cases be satisfied irrespective of the service discipline.

The condition (1.10) will imply the positive Harris recurrence of $X(\cdot)$ provided that the states in the state space S communicate with one another in an appropriate sense. Petite sets are typically employed for this purpose; they will be defined in Section 2. A petite set A has the property that each measurable set B is “equally accessible” from all points in A with respect to a given nontrivial measure.

THEOREM 1.2. *Suppose that a JSQ network is subcritical, satisfies (1.7), and that $A_M = \{x : \|x\| \leq M\}$ is petite for each $M > 0$ for the norm in (1.8). Then $X(\cdot)$ is positive Harris recurrent.*

Theorem 1.2 will follow from Theorem 1.1 by standard reasoning. More detail is provided in Section 2.

A standard criterion that ensures the above sets A_M are petite is given by the following two conditions on the interarrival times. In various works on stability (e.g., [3, 8, 11]), these conditions are employed rather than the more abstract notion of petite set. The first condition is that the distribution $G_k(\cdot)$ is unbounded for each k , that is,

$$(1.11) \quad \bar{G}_k(y) \stackrel{\text{def}}{=} 1 - G_k(y) > 0 \quad \text{for all } y.$$

The second condition is that, for some $\ell_k \in \mathbb{Z}_+$, the ℓ_k -fold convolution $G_k^{*\ell_k}(\cdot)$ of $G_k(\cdot)$ and Lebesgue measure are not mutually singular. That is, for some nonnegative $q_k(\cdot)$ with $\int_0^\infty q_k(s) ds > 0$,

$$(1.12) \quad G_k^{*\ell_k}(d) - G_k^{*\ell_k}(c) \geq \int_c^d q_k(s) ds$$

for all $c < d$. When the interarrival times are exponentially distributed, both (1.11) and (1.12) are immediate. More detail is given in Section 2.

We therefore have the following corollary of Theorem 1.2. As noted earlier, (1.7) is automatic in this setting.

COROLLARY 1.1. *Suppose that a subcritical JSQ network has interarrival times that satisfy (1.11) and (1.12). Then $X(\cdot)$ is positive Harris recurrent.*

By employing Theorem 1.2 and the bounds obtained in the derivation of Theorem 1.1, one can obtain uniform bounds on the equilibria restricted to individual queues and to individual arrival streams for families of JSQ networks. Such a family \mathcal{A} will be required to satisfy the following uniformity conditions on the service and interarrival distributions $F_j^{(a)}(\cdot)$ and $G_k^{(a)}(\cdot)$, for $a \in \mathcal{A}$:

$$(1.13) \quad \sup_{a \in \mathcal{A}} \max_j \mu_j^{(a)} \int_{M/\mu_j^{(a)}}^\infty y F_j^{(a)}(dy) \rightarrow 0 \quad \text{as } M \rightarrow \infty$$

and

$$(1.14) \quad \sup_{a \in \mathcal{A}} \max_k \alpha_k^{(a)} \int_{M/\alpha_k^{(a)}}^\infty y G_k^{(a)}(dy) \rightarrow 0 \quad \text{as } M \rightarrow \infty.$$

We require

$$(1.15) \quad \sup_{a \in \mathcal{A}} \rho_i^{(a)} < 1,$$

where $\rho_i^{(a)} = \rho_1^{(a)}$ or $\rho_i^{(a)} = \rho_2^{(a)}$, depending on whether the network is class independent or station independent. Setting

$$(\overset{\circ}{m}^{\text{ratio}})^{(a)} = \begin{cases} 1 & \text{for class independent networks,} \\ \frac{\max_{k,B} m_{k,B}^{(a)}}{\min_{k,B} m_{k,B}^{(a)}} & \text{for station independent networks,} \end{cases}$$

we also require that

$$(1.16) \quad \overset{\circ}{m}^{\text{ratio}} \stackrel{\text{def}}{=} \sup_{a \in \mathcal{A}} (\overset{\circ}{m}^{\text{ratio}})^{(a)} < \infty.$$

We also need a uniform version of (1.7), namely that, for some $\Gamma \in \mathbb{Z}_{+,0}$ and $h(\cdot)$, with $h(t) > 0$ for all t ,

$$(1.17) \quad P_x^{(a)}(\text{at most } \Gamma \text{ potential arrivals at } n \text{ occur over } (0, (m^{\max})^{(a)}t]) \geq h(t)$$

for all $a \in \mathcal{A}$, x , t and n . Here, $P_x^{(a)}(\cdot)$ denotes the transition kernel with respect to a . When the selection rules $p_{k,B}^{(a)}$ satisfy $p_{k,B}^{(a)} = 0$ for $|B| \geq M$ for some M not depending on a , we will say the selection rules have *uniformly bounded support*.

In many cases of interest, conditions (1.13)–(1.17) are not difficult to check. For instance, when each network has a single Poisson arrival stream and $F_j^{(a)}(\cdot)$ does not depend on j or a , all conditions except for (1.15) and (1.17) are automatic. When, in addition, the selection rules have uniformly bounded support, then (1.17) also holds.

By Theorem 1.2, for any subcritical JSQ network $a \in \mathcal{A}$, with A_M petite for $M > 0$, the underlying Markov process $X^{(a)}(\cdot)$ is positive Harris recurrent; this follows by employing the analog of (1.8),

$$(1.18) \quad \|x\|^{(a)} = \|x\|_L^{(a)} + \|x\|_R^{(a)} + \|x\|_A^{(a)}.$$

In the statement of Theorem 1.3, we also employ the *local norm* at $n = 1, \dots, N^{(a)}$,

$$(1.19) \quad |x|_n^{(a)} = z_n + \ell_n^{(a)} + w_n^{(a)} \quad \text{for } x \in \mathcal{S}^{(a)}.$$

Here, z_n is the number of jobs at queue n and $w_n^{(a)}$ is the *weighted workload* at n , that is,

$$(1.20) \quad w_n^{(a)} = \sum_{i=1}^{z_n} w_{n,i}^{(a)} = \sum_{i=1}^{z_n} \mu_{j_{n,i}}^{(a)} v_{n,i},$$

where $\mu_{j_{n,i}}^{(a)}$ denotes the service rate of the i th job at queue n and $v_{n,i}$ is the residual service time of this job. The term $\ell_n^{(a)}$ is the *maximum weighted age* at n , that is,

$$(1.21) \quad \ell_n^{(a)} = \max_{i=1, \dots, z_n} \ell_{n,i}^{(a)} = \max_{i=1, \dots, z_n} \mu_{j_{n,i}}^{(a)} o_{n,i},$$

where $o_{n,i}$ is the age of the i th job at queue n (the time since its arrival at n). We set

$$(1.22) \quad s_k^{(a)} = \alpha_k^{(a)} u_k,$$

where u_k is the residual interarrival time at the arrival stream k ; we refer to $s_k^{(a)}$ as the *weighted interarrival time*. (Since x and z do not explicitly depend on $a \in \mathcal{A}$, the corresponding superscripts are omitted.) We will employ, in Section 4, some of the terminology and conditions introduced in (1.13)–(1.22), when demonstrating Theorem 1.1 for a given JSQ network.

Since $X^{(a)}(\cdot)$ is positive Harris recurrent, it will have a unique equilibrium measure on $S^{(a)}$, which we denote by $\mathcal{E}^{(a)}$. In order to ensure that the marginal distribution at a given queue n does not depend on n , we also assume that, for some subgroup $H^{(a)}$ of the permutation group on $B_{N^{(a)}+K^{(a)}}$, with queues being mapped to queues, arrival streams to arrival streams and on which all queues (but not necessarily all arrival streams) communicate,

$$(1.23) \quad X_\pi^{(a)}(\cdot) \text{ and } X^{(a)}(\cdot) \text{ are stochastically equivalent for all } \pi \in H^{(a)}.$$

[That is, $X_\pi^{(a)}(\cdot)$ and $X^{(a)}(\cdot)$ have the same joint distributions.] Here, $X_\pi^{(a)}(\cdot)$ is the stochastic process induced from $X^{(a)}(\cdot)$ by permuting the queues and arrival streams according to π . We will call such a JSQ network *symmetric*. As an example of such a JSQ network, one can consider N queues arranged uniformly along a circle, along with N arrival streams that alternate with the queues, with jobs arriving at a given stream by selecting the shortest queue within a preassigned distance of the stream, and the service discipline and service and interarrival distributions not depending on the queue, arrival stream or selection set.

Employing the preceding conditions, we now state our third main result. Here and later on, $\mathcal{E}^{(a)}(\cdot)$ denotes the probability of the indicated event with respect to the measure $\mathcal{E}^{(a)}$, and X and $S^{(a)}$ denote the random variables corresponding to x and $s^{(a)}$. [The random variables $S^{(a)} = (S_1^{(a)}, \dots, S_N^{(a)})$ should not be confused with the state space.]

THEOREM 1.3. *Suppose that a family \mathcal{A} of JSQ networks satisfies the uniformity conditions (1.13)–(1.17) and (1.23), and that, for each network $a \in \mathcal{A}$, $A_M = \{x : \|x\|^{(a)} \leq M\}$ is petite for each $M > 0$ with respect to the norms in (1.18). Then*

$$(1.24) \quad \sup_{a \in \mathcal{A}} \mathcal{E}^{(a)}(|X|_n^{(a)} > M) \rightarrow 0 \quad \text{as } M \rightarrow \infty,$$

for each n , and

$$(1.25) \quad \sup_{a \in \mathcal{A}} \max_k \mathcal{E}^{(a)}(S_k^{(a)} > M) \rightarrow 0 \quad \text{as } M \rightarrow \infty.$$

The limits (1.24) and (1.25) supply uniform bounds on the number of jobs, maximum weighted age and weighted workload at each individual queue n , and on the weighted interarrival times at each arrival stream k . When these limits hold, we say the equilibria $\mathcal{E}^{(a)}$, for $a \in \mathcal{A}$, are *locally bounded*. Note that, on account of (1.23), the probabilities in (1.24) do not depend on n .

In many cases of interest, the conditions in Theorem 1.3 are not difficult to check. As mentioned in conjunction with (1.13)–(1.17), when each member of \mathcal{A} has a single Poisson arrival stream, $F_j^{(a)}(\cdot)$ does not depend on j or a , and the selection rules have uniformly bounded support, then all of these properties except (1.15) hold. Since (1.11) and (1.12) hold for Poisson arrivals, the sets A_M are petite. Also, under (1.23), the traffic intensity can be written as in (1.6). We therefore obtain the following corollary of Theorem 1.3. Recall that a selection rule is mean field if, for a given D , each nonrepeating D -tuple is chosen with equal probability from among the $N^{(a)}$ queues, and note that JSQ networks with a mean field selection rule and for which $F_j^{(a)}$ does not depend on j or a are also symmetric.

COROLLARY 1.2. *Suppose that each member of a family \mathcal{A} of JSQ networks has a single Poisson arrival stream, that $F_j^{(a)}$ does not depend on j or a , that the selection rules have uniformly bounded support, and that*

$$(1.26) \quad \sup_{a \in \mathcal{A}} \alpha^{(a)} m / N^{(a)} < 1.$$

If (1.23) is satisfied, then (1.24) and (1.25) hold. In particular, if the selection rules are mean field, then (1.24) and (1.25) hold.

As mentioned earlier, when a JSQ network has a single Poisson arrival stream, one can omit the interarrival times from the state space descriptor. In this case, the limit (1.25) is no longer relevant in Theorem 1.3 and Corollary 1.2.

When \mathcal{A} is given by a family of networks indexed by the number N of queues, Theorem 1.3 provides local bounds on $\mathcal{E}^{(N)}$ as $N \rightarrow \infty$. These bounds can be used to show the relative compactness of the restriction of $\mathcal{E}^{(N)}$ to finite sets of queues; this is done in Theorem 6.1. As mentioned earlier, these local bounds and relative compactness of the sequence provide a framework for approximating the corresponding marginal distributions for large N [6, 7]. In this context, one employs appropriate mean field equations corresponding to the marginal distributions of the equilibrium $\mathcal{E}^{(\infty)}$ of a limiting infinite system. Under appropriate assumptions on the service times, the solutions of these mean field equations satisfy bounds on the number of jobs that are on the same order as in the limit in (1.1).

In Section 7, we present a family of mean field, symmetric networks, with a single Poisson arrival stream, $N = D = 2$, and an appropriate service discipline that illustrates how the JSQ rule can produce equilibria for which the typical workload is incredibly large, much larger than the workload for the analogous network with $D = 1$. So, in terms of workload, the JSQ rule can sometimes yield poor results.

Outline of the paper and main ideas. In Section 2, we will provide a brief background on Markov processes that will be relevant to the space S and Markov process $X(\cdot)$ employed in the introduction, and we will provide a more detailed construction of S and $X(\cdot)$. We will also provide the background that is needed to derive Theorem 1.2 from Theorem 1.1 and to obtain Corollary 1.1. The machinery for this is standard in the context of queueing networks and is easily modified so as to apply to JSQ networks.

In Section 3, we provide an alternative formulation of the traffic intensities in (1.2) and (1.3) that we employ in successive sections. This formulation will enable us to compare JSQ networks to networks with appropriate random assignment of jobs to queues.

Theorem 1.1 is demonstrated in Sections 4 and 5. The main tool is an appropriate Lyapunov function that is given in terms of the norms $\|x\|_L$, $\|x\|_R$ and $\|x\|_A$ in (1.8). Our analysis involves decomposing time into random intervals over which no jobs enter the network. Over each such interval, the evolution of the system is deterministic and $\|X(t)\|$ is shown to be decreasing at large values. At times t where a job enters the network, the average value of $\|X(t)\| - \|X(t-)\|$ is shown to be negative. Applying the strong Markov property and iterating over such intervals until time τ_M , $\tau_M = \inf\{t : \|X(t)\| \leq M\}$, will imply (1.10) of Theorem 1.1. At the end of Section 5, we briefly discuss the networks mentioned at the beginning of the introduction where jobs are assigned to the queue with the least remaining work. We refer to such networks as *join the least loaded queue* (JLLQ) networks. The JLLQ rule is easier to handle than the JSQ rule. It is analyzed in [11] using fluid limits; here, we mention an alternative approach. At the end of the section, we also mention analogs of Theorem 1.1 where the state space S is modified.

Theorem 1.3 is demonstrated in Section 6. The basic idea there is to employ estimates from Sections 4 and 5 to obtain lower bounds on the rate at which $\|X^{(a)}(t)\|^{(a)}$ decreases at large values of $|X^{(a)}(t)|_n^{(a)}$. If these values occur with too high a probability at a given $a \in \mathcal{A}$, it will follow that, for appropriate t ,

$$E_{\mathcal{E}^{(a)}}[\|X^{(a)}(t)\|^{(a)}] < E_{\mathcal{E}^{(a)}}[\|X^{(a)}(0)\|^{(a)}].$$

Since $\mathcal{E}^{(a)}$ is an equilibrium, this is not possible, which gives us upper bounds on the probabilities in (1.24) and (1.25). Theorem 6.1 follows quickly from Theorem 1.3 and Prohorov's theorem.

In Section 7, we discuss the family of networks mentioned earlier, with $N = D = 2$, whose equilibria have very large workload. This behavior arises because of the manner in which the service discipline restricts service for jobs with large residual service time. The main result is given in Theorem 7.1. Because of the length of the argument, we provide a condensed proof of part of the theorem.

Notation. For the reader's convenience, we mention here some of the notation in the paper. The term x indicates a state in the state space S and the corresponding

term $X(t)$ (or X) indicates a random state at time t (or with respect to a given measure); quantities such as z_n and $Z_n(t)$, and $w_{n,i}$ and $W_{n,i}(t)$ (or, Z_n and $W_{n,i}$) play analogous roles. The results in this paper are for class independent and for station independent networks; in order to avoid repeating all statements and proofs for the two cases, we employ notation with the symbol \circ , such as $\mathring{m}_{k,A}$, which will have different meanings in the two cases. The main norms with which we will work are $\|\cdot\|$, $\|\cdot\|_L$, $\|\cdot\|_R$ and $\|\cdot\|_A$. The symbols L , R and A also appear in other contexts, such as for the numbers L_i , $i = 1, \dots, 4$, service effort per job $R_{n,i}$ and sets A ; the meaning should always be clear from the context. The symbols \mathbb{Z}_+ and \mathbb{R}_+ denote the positive integers and positive real numbers, with $\mathbb{Z}_{+,0} = \mathbb{Z}_+ \cup \{0\}$ and $\mathbb{R}_{+,0} = \mathbb{R}_+ \cup \{0\}$; $\lfloor y \rfloor$ denotes the integer part of $y \in \mathbb{R}_+$ and $\mathbf{1}\{A\}$ denotes the indicator function of the event A .

2. Markov process background. In this section, we provide a more detailed description of the construction of the Markov process $X(\cdot)$ that underlies a JSQ network. We then show how Theorem 1.2 and its corollary follow from Theorem 1.1. Analogs of this material for queueing networks are given in Bramson [4], and, for networks with weighted max–min fair policies, in Bramson [5]. Because of the similarity of these settings, we present a summary here and refer the reader to [4] for more detail.

Construction of the Markov process. We define the state space S to be the set

$$(2.1) \quad (\mathbb{Z}^3 \times 2^{B_N} \times \mathbb{R}^3)^\infty \times \mathbb{R}^K$$

subject to the following constraints. The components s_k , $k = 1, \dots, K$, of \mathbb{R}^K are all positive; they correspond to the residual interarrival times u_k of the K arrival streams, scaled by the arrival rates α_k as in (1.22). Only a finite number of the 7-tuples of coordinates of $(\mathbb{Z}^3 \times 2^{B_N} \times \mathbb{R}^3)^\infty$ are nonzero. Such a 7-tuple corresponds to a particular job in the network: the first coordinate n , $n = 1, \dots, N$, corresponds to the queue of the job and the second coordinate i , $i = 1, \dots, z_n$, gives its rank at the queue based on the time of arrival there, with “older” jobs receiving a lower rank. The third and fourth coordinates k , $k = 1, \dots, K$, and A , $\emptyset \neq A \subseteq B_N$, correspond to the arrival stream and selection set of the job when it entered its queue. The fifth coordinate ℓ , $\ell \geq 0$, measures the age o of the job, scaled by $\mu_{j_n,i}$, as in (1.21); the sixth coordinate w , $w > 0$, measures its residual service time v , scaled by $\mu_{j_n,i}$, as in (1.20); and the last coordinate r , $r \in [0, 1]$, is the current service effort devoted to the job. (If one wishes, one can introduce other coordinates in the state space descriptor, such as the elapsed time since the last arrival from each stream, or the amount of service already received by each job.) Since the discipline is assumed to be nonidling, the sum of the last coordinates for all jobs at a given nonempty queue must equal 1. Note that for given N and K , the state space S constructed in this manner is unique.

The last five coordinates may be considered as functions of the first two, and written as $k_{n,i}$, $A_{n,i}$, $\ell_{n,i}$, $w_{n,i}$ and $r_{n,i}$. The third and fourth coordinates, $k_{n,i}$ and $A_{n,i}$, are needed because of how $\|\cdot\|$ is defined in Section 4 and can be omitted for class independent networks. Various coordinates can also be omitted for particular service disciplines (such as FIFO).

For given $N' \leq N$ and $K' = K$, one can define S' as above, but with $n = 1, \dots, N'$. Then S' is the *projection* of S obtained by restricting nonzero 7-tuples to the first N' queues. For $x \in S$, the projection $x' \in S'$ of x is the element obtained by omitting 7-tuples with $n > N'$. One can define projections of S onto spaces S' corresponding to other subsets of $\{1, \dots, N\}$ analogously, but we will not use these in the paper.

Employing the above notation, we construct a metric $d(\cdot, \cdot)$ on S : for given $x, x' \in S$, with the coordinates labelled correspondingly, we set

$$\begin{aligned}
 d(x, x') &= \sum_{n=1}^N \sum_{i=1}^{\infty} ((|\ell_{n,i} - \ell'_{n,i}| + |w_{n,i} - w'_{n,i}| + |r_{n,i} - r'_{n,i}|) \wedge 1) \\
 (2.2) \quad &+ \sum_{n=1}^N \sum_{i=1}^{\infty} (\mathbf{1}\{k_{n,i} \neq k'_{n,i}\} + \mathbf{1}\{A_{n,i} \neq A'_{n,i}\}) \\
 &+ \sum_{n=1}^N |z_n - z'_n| + \sum_{k=1}^K |s_k - s'_k|.
 \end{aligned}$$

One can check that $d(\cdot, \cdot)$ is separable and locally compact; more detail is given on page 82 of [4]. One can also check that the sets $E_M \subset S$, $M > 1$, defined by $z_n \leq M$, $\ell_{n,i} \leq M$, $1/M \leq w_{n,i} \leq M$ and $1/M \leq s_k \leq M$, for all $n = 1, \dots, N$, $k = 1, \dots, K$ and $i = 1, \dots, z_n$ are compact with respect to $d(\cdot, \cdot)$. We equip S with the standard Borel σ -algebra inherited from $d(\cdot, \cdot)$, which we denote by \mathcal{S} . In Lemma 4.1, we will show $\|\cdot\|_L$, $\|\cdot\|_R$ and $\|\cdot\|_A$ are continuous in $d(\cdot, \cdot)$.

At the end of Section 6, we will employ the partial completion \bar{S} of S that is obtained by allowing the weighted residual service times $w_{n,i}$ to take values in $[0, \infty)$ rather than just $(0, \infty)$. Otherwise, the construction of \bar{S} is the same as that just given for S . The metric $\bar{d}(\cdot, \cdot)$ is defined analogously to $d(\cdot, \cdot)$, and the Borel σ -algebra $\bar{\mathcal{S}}$ is defined correspondingly. Under $\bar{d}(\cdot, \cdot)$, the sets $\bar{E}_M \subset \bar{S}$, $M > 1$, defined by $z_n \leq M$, $\ell_{n,i} \leq M$, $0 \leq w_{n,i} \leq M$ and $1/M \leq s_k \leq M$, for all $n = 1, \dots, N$, $k = 1, \dots, K$ and $i = 1, \dots, z_n$ are compact. One can define projections from \bar{S} onto spaces \bar{S}' in the same manner as was done from S onto S' .

The Markov process $X(t)$, $t \geq 0$, underlying the network is defined to be the right continuous process taking values x in S whose evolution is determined by the given JSQ rule together with the assigned service discipline. Jobs (n, i) are allocated service according to rates $R_{n,i}(t)$ (the service effort per job) that are assumed to be constant in between arrivals and departures of jobs at the queues. Over such an interval, $\mathcal{L}_{n,i}(t)$ increases at rate $\mu_{j_{n,i}}$, and $W_{n,i}(t)$ and $S_k(t)$ decrease at

rates $\mu_{j_n,i} R_{n,i}(t)$ and α_k , respectively. [We write $\mathcal{L}_{n,i}(t)$ for the age functions to avoid possible confusion later with constants L_i that will be introduced.] Upon an arrival or departure, rates are re-assigned according to the discipline. The restriction made here for the discipline, that service rates remain constant between arrivals and departures of jobs, is for convenience, and allows one to inductively construct $X(\cdot)$ over increasing times in a simple way. The standard service disciplines satisfy this property. We also note that the construction here is not restricted to the JSQ rule, and applies to other rules for assigning arriving jobs to a queue.

Along the lines of page 85 of [4], a filtration (\mathcal{F}_t) , $t \in [0, \infty]$, can be assigned to $X(\cdot)$ so that $X(\cdot)$ is Borel right and, in particular, is strong Markov. The processes $X(\cdot)$ fall into the class of piecewise-deterministic Markov processes, for which the reader is referred to Davis [9] for more detail.

Recurrence. The Markov process $X(\cdot)$ is said to be *Harris recurrent* if, for some nontrivial σ -finite measure φ ,

$$\varphi(B) > 0 \quad \text{implies} \quad P_x(\eta_B = \infty) = 1 \quad \text{for all } x \in S,$$

where $\eta_B = \int_0^\infty \mathbf{1}\{X(t) \in B\} dt$. If $X(\cdot)$ is Harris recurrent, it possesses a stationary measure π that is unique up to a constant multiple. When π is finite, $X(\cdot)$ is said to be *positive Harris recurrent*.

A practical condition for determining positive Harris recurrence can be given by using petite sets. A nonempty set $A \in \mathcal{S}$ is said to be *petite* if for some fixed probability measure a on $(0, \infty)$ and some nontrivial measure ν on (S, \mathcal{S}) ,

$$\nu(B) \leq \int_0^\infty P^t(x, B)a(dt)$$

for all $x \in A$ and $B \in \mathcal{S}$. Here, $P^t(\cdot, \cdot)$, $t \geq 0$, is the semigroup associated with $X(\cdot)$. As mentioned in the introduction, a petite set A has the property that each set B is “equally accessible” from all points $x \in A$ with respect to the measure ν . Note that any nonempty measurable subset of a petite set is also petite.

For given $\delta > 0$, set

$$\tau^A(\delta) = \inf \{t \geq \delta : X(t) \in A\}$$

and $\tau^A = \tau^A(0)$. Then $\tau^A(\delta)$ is a stopping time. Employing petite sets and $\tau^A(\delta)$, one has the following characterization of Harris recurrence and positive Harris recurrence. (The Markov process and state space need to satisfy minimal regularity conditions, as on page 86 of [4].) The criteria are from Meyn and Tweedie [15]; discrete time analogs of the different parts of the proposition have long been known, see, for instance, Nummelin [17] and Orey [18].

THEOREM 2.1. (a) *A Markov process $X(\cdot)$ is Harris recurrent if and only if there exists a closed petite set A with*

$$(2.3) \quad P_x(\tau^A < \infty) = 1 \quad \text{for all } x \in S.$$

(b) *Suppose the Markov process $X(\cdot)$ is Harris recurrent. Then $X(\cdot)$ is positive Harris recurrent if and only if there exists a closed petite set A such that, for some $\delta > 0$,*

$$(2.4) \quad \sup_{x \in A} E_x[\tau^A(\delta)] < \infty.$$

Theorem 1.2 and its corollary. Theorem 1.2 follows from Theorems 1.1, 2.1 and the elementary continuity result, Lemma 4.1. To see this, note that both conditions (2.3) and (2.4) of Theorem 2.1 are immediate consequences of (1.10) of Theorem 1.1, with $A = A_M$, for appropriate M , and $\delta = 1$ since, in Theorem 1.2, A_M is assumed to be petite. By Lemma 4.1, the norm $\|\cdot\|$ in (1.8) is continuous in the metric $d(\cdot)$, and hence A_M is also closed. It therefore follows from Theorem 2.1 that $X(\cdot)$ is positive Harris recurrent, which implies Theorem 1.2.

Corollary 1.1 follows immediately from Theorem 1.2 and the assertion, before the statement of the corollary, that the sets A_M are petite under the assumptions (1.11) and (1.12). A somewhat stronger version of the analogous assertion for queueing networks is demonstrated in Proposition 4.7 of [4]. (The proposition states that the sets A are uniformly small.) The reasoning is the same in both cases and does not involve the JSQ rule or the service discipline. The argument, in essence, requires that one wait long enough for the network to have at least a given positive probability of being empty; the time t does not depend on x for $\|x\| \leq M$. This will follow from (1.11) and the definition of $\|\cdot\|$ in Section 4, since the work in the network is bounded by a linear function of M . Since the residual interarrival times are also bounded by a linear function of M , by using (1.12), one can also show that the joint distribution function of the residual interarrival times has an absolutely continuous component at this time, whose density is bounded away from 0. It will follow that the set A_M is petite with respect to ν , with a chosen as the point mass at t , if ν is concentrated on the empty states, where it is a small enough multiple of $|\mathcal{R}|$ -dimensional Lebesgue measure restricted to a small cube.

3. A useful routing lemma. In this section, we rephrase the conditions (1.2) and (1.3) that are used to define subcriticality for class independent and station independent JSQ networks by using sums that will be more convenient for us to work with when proving Theorem 1.1 in Sections 4 and 5. For this, we employ Lemma 3.1. The desired sums for class independent and station independent JSQ networks are then given in the following corollary.

For the lemma, we employ the following notation. We consider $\beta_{k,A,n} \geq 0$ for $k = 1, \dots, K, n = 1, \dots, N$ and $A \subseteq B_N$, where $B_N = \{1, \dots, N\}$, and assume that, for each k and A ,

$$(3.1) \quad \beta_{k,A,n} > 0 \quad \text{iff} \quad \beta_{k,A,n'} > 0 \quad \text{for all } n, n' \in A,$$

that is, whether or not $\beta_{k,A,n}$ is zero does not depend on n , for $n \in A$. For each $B \subseteq B_N$, let $\{r_{B,n}, n = 1, \dots, N\}$ be a probability distribution concentrated on B , with $r_{B,n} > 0$ for $n \in B$, so that for each k, B , and $A \subseteq B$, and n restricted to B ,

$$(3.2) \quad \gamma_{k,A,B} \stackrel{\text{def}}{=} r_{B,n} \beta_{k,A,n} \quad \text{does not depend on } n.$$

LEMMA 3.1. *Suppose that $\beta_{k,A,n}$, $\gamma_{k,A,B}$ and $r_{B,n}$ are chosen as in (3.1)–(3.2), and moreover that $\gamma_{k,A,B}$ satisfies*

$$(3.3) \quad \sum_k \sum_{A \subseteq B} \gamma_{k,A,B} \leq \rho \quad \text{for all } B \subseteq B_N$$

for some ρ . Then, for each k and A , there is a probability distribution $\{q_{k,A,n}, n = 1, \dots, N\}$ concentrated on A , so that

$$(3.4) \quad \sum_k \sum_{A \subseteq B_N} \beta_{k,A,n} q_{k,A,n} \leq \rho \quad \text{for all } n.$$

As a consequence of Lemma 3.1, we obtain the following two inequalities from (1.2) and (1.3).

COROLLARY 3.1. (a) *Suppose that (1.2) holds for a class independent JSQ network. Then, for each k and A , there is a probability distribution $q_{k,A,n}$ concentrated on A so that*

$$(3.5) \quad \sum_k \sum_{A \subseteq B_N} \alpha_k p_{k,A} q_{k,A,n} m_n \leq \rho_1 \quad \text{for all } n.$$

(b) *Suppose that (1.3) holds for a station independent JSQ network. Then, for each k and A , there is a probability distribution $q_{k,A,n}$ concentrated on A so that*

$$(3.6) \quad \sum_k \sum_{A \subseteq B_N} \alpha_k p_{k,A} q_{k,A,n} m_{k,A} \leq \rho_2 \quad \text{for all } n.$$

Corollary 3.1 converts (1.2) and (1.3) into (3.5) and (3.6), which will be easier to work with in Section 4. When $\rho_1 < 1$, respectively, $\rho_2 < 1$, these inequalities imply that the network, where incoming jobs are assigned the queue n with probabilities $q_{k,A,n}$, is subcritical. On the other hand, by substituting B for B_N in the inner sum and summing over n , one can check that the inequalities in (3.5), respectively, (3.6), cannot be strict for all n . These conditions therefore give an alternative characterization for subcriticality when $\rho_i < 1$. (We will not need this in the paper.)

One can check that when the JSQ network is symmetric, one may set $q_{k,A,n} = 1/|A|$ for $n \in A$. This follows by summing the left-hand side of (3.6) over all n and comparing the sum with the right-hand side of (1.3), for $B = B_N$.

PROOF OF COROLLARY 3.1. (a) We apply Lemma 3.1, setting

$$(3.7) \quad \beta_{k,A,n} = \alpha_k p_{k,A} m_n, \quad r_{B,n} = \mu_n \mathbf{1}\{n \in B\} / \sum_{n' \in B} \mu_{n'}, \quad \rho = \rho_1.$$

It is easy to check that the conditions (3.1) and (3.2) are satisfied with this choice of $\beta_{k,A,n}$ and $r_{B,n}$. Substitution of these quantities into (3.3) gives the quantity in braces on the right-hand side of (1.2) for each choice of B , and substitution into (3.4) gives (3.5). Part (a) of the corollary therefore follows from the lemma.

(b) The argument for this part is analogous to the first part; here, we set

$$(3.8) \quad \beta_{k,A,n} = \alpha_k p_{k,A} m_{k,A}, \quad r_{B,n} = \mathbf{1}\{n \in B\} / |B|, \quad \rho = \rho_2.$$

It is easy to check that the conditions (3.1)–(3.2) are again satisfied. Substitution into (3.3) gives the quantity in braces on the right-hand side of (1.3) and substitution into (3.4) gives (3.6). Part (b) of the corollary therefore also follows from the lemma. \square

We now prove the lemma.

PROOF OF LEMMA 3.1. For a family of probability distributions $q_{k,A,n}$, indexed by k and A , and concentrated on A , set $\zeta_{k,A,n}^{(q)} = \beta_{k,A,n} q_{k,A,n}$,

$$(3.9) \quad V(q) = \sum_n \left(\sum_k \sum_{A \subseteq B_N} \zeta_{k,A,n}^{(q)} - \rho \right)^2$$

and

$$(3.10) \quad V^{\min} = \min_q V(q).$$

One can check that $V(q)$ is continuous in q and that the set of q is compact with respect to the metric

$$(3.11) \quad d(q, q') = \max_{k,A,n} |q_{k,A,n} - q'_{k,A,n}|.$$

So, V^{\min} is attained at some q^{\min} . We set

$$(3.12) \quad \begin{aligned} A^{\min} &= \left\{ n : \sum_k \sum_{A \subseteq B_N} \zeta_{k,A,n}^{\min} > \rho \right\}, \\ \bar{A}^{\min} &= \left\{ n : \sum_k \sum_{A \subseteq B_N} \zeta_{k,A,n}^{\min} \geq \rho \right\}, \end{aligned}$$

where $\zeta_{k,A,n}^{\min} \stackrel{\text{def}}{=} \beta_{k,A,n} q_{k,A,n}^{\min}$. In order to show (3.4), it suffices to show $A^{\min} = \emptyset$.

We first claim that for any $k', A' \subseteq B_N$ and $n_1 \in \bar{A}^{\min}$,

$$(3.13) \quad A' \subseteq \bar{A}^{\min} \quad \text{if } \zeta_{k',A',n_1}^{\min} > 0.$$

We argue by contradiction and show that if (3.13) is violated for some k', A' and n_1 , then for appropriate \tilde{q} , $V(\tilde{q}) < V(q^{\min})$, which is not possible. The proof of this does not use (3.3).

For such A' and n_1 , $n_1 \in A'$ since $\zeta_{k',A'}^{\min}$ is concentrated on A' . We choose $n_2 \in A' - \bar{A}^{\min}$ and define a new family of probability distributions $\tilde{q}_{k,A}$ by

$$\begin{aligned} \tilde{q}_{k,A,n} &= q_{k,A,n}^{\min} && \text{unless } k = k', A = A' \text{ and either } n = n_1 \text{ or } n = n_2, \\ \tilde{q}_{k',A',n_1} &= q_{k',A',n_1}^{\min} - \varepsilon, \\ \tilde{q}_{k',A',n_2} &= q_{k',A',n_2}^{\min} + \varepsilon, \end{aligned}$$

where $\varepsilon > 0$ is small.

For small enough $\varepsilon > 0$,

$$\begin{aligned} (3.14) \quad & \left(\sum_k \sum_{A \subseteq B_N} \tilde{\zeta}_{k,A,n_1} - \rho \right)^2 - \left(\sum_k \sum_{A \subseteq B_N} \zeta_{k,A,n_1}^{\min} - \rho \right)^2 \\ & < \left(\sum_k \sum_{A \subseteq B_N} \zeta_{k,A,n_2}^{\min} - \rho \right)^2 - \left(\sum_k \sum_{A \subseteq B_N} \tilde{\zeta}_{k,A,n_2} - \rho \right)^2, \end{aligned}$$

where $\tilde{\zeta}_{k,A,n} = \beta_{k,A,n} \tilde{q}_{k,A,n}$. The inequality is obvious when $n_1 \in A^{\min}$, since the left-hand side of (3.14) will be negative and the right-hand side will be positive. When $n_1 \in \bar{A}^{\min} - A^{\min}$, the inequality is also true since the left-hand side is bounded above by a multiple of ε^2 (with the second term being 0), and the right-hand side is bounded below by a multiple of ε . For $n \neq n_1, n_2$,

$$(3.15) \quad \left(\sum_k \sum_{A \subseteq B_N} \tilde{\zeta}_{k,A,n} - \rho \right)^2 = \left(\sum_k \sum_{A \subseteq B_N} \zeta_{k,A,n}^{\min} - \rho \right)^2.$$

So, by (3.14) and (3.15), $V(\tilde{q}) < V(q^{\min})$, which contradicts the definition of q^{\min} . This shows (3.13).

Employing (3.13), we now show that $A^{\min} = \emptyset$. One has

$$\begin{aligned} (3.16) \quad \rho &\geq \sum_k \sum_{A \subseteq \bar{A}^{\min}} \gamma_{k,A,\bar{A}^{\min}} = \sum_k \sum_{A \subseteq \bar{A}^{\min}} \sum_{n=1}^N \gamma_{k,A,\bar{A}^{\min}} q_{k,A,n}^{\min} \\ &= \sum_k \sum_{A \subseteq \bar{A}^{\min}} \sum_{n \in \bar{A}^{\min}} r_{\bar{A}^{\min},n} \beta_{k,A,n} q_{k,A,n}^{\min} \\ &= \sum_{n \in \bar{A}^{\min}} r_{\bar{A}^{\min},n} \sum_k \sum_{A \subseteq B_N} \zeta_{k,A,n}^{\min}. \end{aligned}$$

The inequality follows from (3.3), the second equality follows from (3.2) since $r_{B,\cdot}$ is concentrated on B , and the third equality follows from the definition of $\zeta_{k,A,n}^{\min}$

and (3.13). But, on account of the definitions of A^{\min} and \bar{A}^{\min} , the last quantity in (3.16) is at least

$$\rho \sum_{n \in \bar{A}^{\min}} r_{A^{\min},n} = \rho,$$

with strict inequality holding if $A^{\min} \neq \emptyset$ because all the terms in the sum are strictly positive. Since the strict inequality contradicts (3.16), this implies $A^{\min} = \emptyset$, and hence the inequality in (3.4), as desired. \square

4. Definitions of norms and basic inequalities. This section introduces the norms and provides the basic inequalities we will need in Sections 5 and 6 for the proofs of Theorems 1.1 and 1.3. The section consists of two subsections. We first define the norms appearing in (1.8), $\|\cdot\|_L$, $\|\cdot\|_R$ and $\|\cdot\|_A$, in terms of which $\|\cdot\|$ was defined. We then state and prove Propositions 4.1 and 4.2. These propositions give inequalities on the decrease of $\|\cdot\|$ and lie at the heart of the analysis in Sections 5 and 6. Proposition 4.2 is the only place in the first six sections of the paper where the JSQ property is employed.

Definition of the norms. In (1.8), we defined the norm $\|\cdot\|$ in terms of the norms $\|\cdot\|_L$, $\|\cdot\|_R$ and $\|\cdot\|_A$. We now define these norms.

Recall from Section 1 that z_n denotes the number of jobs at queue n ; w_n denotes the weighted workload at the queue, which is defined in (1.20) along with $w_{n,i}$; $r_{n,i}$ denotes the service effort for job (n, i) ; and s_k denotes the weighted interarrival time at the arrival stream k , and is given by (1.22). The notation $Z_n(\cdot)$, $W_n(\cdot)$, $W_{n,i}(\cdot)$, $R_{n,i}(\cdot)$ and $S_k(\cdot)$ will be used for the corresponding quantities of the process $X(\cdot)$. We also employ the arrival rates α_k , mean service times m_j , service rates μ_j and transition probabilities $p_{k,A}$ that were introduced earlier, as well as the transition probabilities $q_{k,A,n}$ that are given in Corollary 3.1 for class independent and station independent JSQ networks. We will set $\varepsilon_1 = 1 - \rho_1$ for class independent networks and $\varepsilon_1 = 1 - \rho_2$ for station independent networks, where ρ_i are the traffic intensities.

For $x \in \mathcal{S}$, set

$$(4.1) \quad \|x\|_L = \sum_{n=1}^N \|x\|_{L,n}, \quad \|x\|_R = \sum_{n=1}^N \|x\|_{R,n}, \quad \|x\|_A = \sum_{k=1}^K \|x\|_{A,k}.$$

(The subscripts L , R and A are mnemonics for “left,” “right” and “arrivals.”) We define these individual components as follows. Since $\|x\|_{L,n}$ is defined in terms of quantities obtained from $\|x\|_{R,n}$, we define the latter first.

For $n = 1, \dots, N$, set

$$(4.2) \quad \|x\|_{R,n} = \sum_{i=1}^{z_n} \overset{\circ}{m}_{k_{n,i},A_{n,i}} \psi_W(w_{n,i}),$$

where $w_{n,i} = \mu_{j_n,i} v_{n,i}$. The other components in (4.2) are defined as follows:

$$(4.3) \quad \mathring{m}_{k,A} = \begin{cases} 1 & \text{for class independent networks,} \\ m_{k,A} & \text{for station independent networks,} \end{cases}$$

and $k_{n,i}$ and $A_{n,i}$ denote the arrival stream and selection set for the i th job currently at queue n . The function $\psi_W : \mathbb{R}_{+,0} \rightarrow \mathbb{R}_{+,0}$ is required to be continuously differentiable, with $\psi_W(0) = 0$, $\psi'_W(y) > 0$, $\psi'_W(y) \nearrow \infty$ as $y \nearrow \infty$, and

$$(4.4) \quad \int_0^\infty \psi_W(\mu_j y) F_j(dy) \leq \varepsilon_2 \quad \text{for all } j,$$

where $\varepsilon_2 = (\varepsilon_1)^2/40$. Since $F_j(\cdot)$ has finite mean, it is not difficult to choose such a $\psi_W(\cdot)$.

The norm $\|\cdot\|_R$ will be the main contributor to $\|\cdot\|$ for jobs with large residual service times; service of such a job (n, i) sharply decreases $\|\cdot\|_R$ when $\psi'_W(w_{n,i})$ is large. The term $\mathring{m}_{k,A}$ is needed because of the different definitions of the traffic intensities ρ_1 and ρ_2 for class independent and station independent networks.

For $n = 1, \dots, N$, we set

$$(4.5) \quad \|x\|_{L,n} = \left(\sum_{i=1}^{z_n} \mathring{m}_{k_{n,i}, A_{n,i}} (w_{n,i}^+ \wedge L_2) \right) \psi_Z(z_n).$$

Here,

$$(4.6) \quad w_{n,i}^+ = w_{n,i} + \varepsilon_2$$

and

$$(4.7) \quad \psi_Z(y) = \begin{cases} \varepsilon_1 + (\varepsilon_3/L_2) \log(y + 1) & \text{for } y \in [0, L_3], \\ \varepsilon_1 + (\varepsilon_3/L_2) \log(L_3 + 1) & \text{for } y > L_3 \end{cases}$$

for a small $\varepsilon_3 > 0$, which will be defined in (4.20) in terms of ε_1 and other quantities.

We choose L_2 in (4.5) and (4.7) so that

$$(4.8) \quad \psi'_W(L_2) = 2L_1$$

for given L_1 , with $L_1 \geq 4$ and L_1 large enough so that $L_2 \geq \varepsilon_2$. We will specify L_1 later as we find convenient. We choose L_3 so that

$$(4.9) \quad \psi_Z(L_3) = \varepsilon_1 + (\varepsilon_3/L_2) \log(L_3 + 1) = L_1;$$

it follows that when $w \geq L_2$,

$$(4.10) \quad \psi_Z(y) \leq \frac{1}{2} \psi'_W(w) \quad \text{for all } y.$$

Note that $L_2 \rightarrow \infty$ and $L_3 \rightarrow \infty$ as $L_1 \rightarrow \infty$. The inequality (4.10) is used in Proposition 4.1, and will tell us that, for large residual service times, the norm $\|\cdot\|_R$ is “more powerful” than $\|\cdot\|_L$.

The following provides some motivation for the definition of $\|x\|_{L,n}$. The norm $\|\cdot\|_L$ will be the main contributor to $\|\cdot\|$ for jobs with moderate residual service times. The terms $w_{n,i}^+ \wedge L_2$ and $\psi_Z(z_n)$ are each bounded, with the term $w_{n,i}^+$ decreasing continuously over time as the corresponding job is served; we employ $w_{n,i}^+$ rather than $w_{n,i}$, in (4.5), to ensure $\|x\|_{L,n} \rightarrow \infty$ as $z_n \rightarrow \infty$. The inclusion of the term $\psi_Z(z_n)$, which is nondecreasing in z_n , is reasonable since more jobs at a queue should correspond to a greater value of $\|\cdot\|_L$.

For $k = 1, \dots, K$, we set

$$(4.11) \quad \|x\|_{A,k} = (1 + \varepsilon_1/2) \sum_{A \subseteq B_N} \left(\sum_{n=1}^N p_{k,A} q_{k,A,n} \dot{m}_{k,A} \psi_Z(z_n) \right) \psi_A(s_k).$$

Here, $q_{k,A,n}$ is chosen as in the corollary to Lemma 3.1 and $s_k = \alpha_k u_k$. We require that $\psi_A(\cdot)$ be locally Lipschitz on $[0, \infty)$, with

$$(4.12) \quad \psi_A(y) = M_1 - y \quad \text{for } y \in [0, M_1],$$

and $\psi'_A(y) > 0$ for $y \in (M_1, \infty)$ and appropriate $M_1 \geq 1$, with $\psi'_A(y) \nearrow \infty$ as $y \nearrow \infty$, so that

$$(4.13) \quad \int_{M_1/\alpha_k}^\infty (\psi_A(\alpha_k y) + \alpha_k y) G_k(dy) \leq \varepsilon_2.$$

Since $G_k(\cdot)$ has finite mean, it is not difficult to choose such M_1 and $\psi_A(\cdot)$. Because $G_k(\cdot)$ has mean $1/\alpha_k$ and (4.12) is satisfied, (4.13) implies that

$$(4.14) \quad M_1 - \int_0^\infty \psi_A(\alpha_k y) G_k(dy) \geq 1 - \varepsilon_2,$$

which will be used in the proof of Proposition 5.1.

The norm $\|\cdot\|_A$ is chosen so that it interfaces properly with $\|\cdot\|_L$: in between arrivals, $\|X(t)\|_A$ will increase more slowly than $\|X(t)\|_L + \|X(t)\|_R$ decreases; at an arrival, the average decrease of $\|X(t)\|_A$ will more than offset the increase in $\|X(t)\|_L + \|X(t)\|_R$ because of (4.4) and (4.14), and the choice of ε_3 in (4.7).

Note that the weighted ages $\ell_{n,i}$ are not employed in the definition of $\|\cdot\|$. They are not needed, in particular, to show petiteness of bounded sets in $\|\cdot\|$ under (1.11) and (1.12), since they do not appear when $z_n = 0$. (See, e.g., the end of Section 2.)

In Section 2, we demonstrated Theorem 1.2 by employing Theorem 1.1 and Lemma 4.1, with the latter asserting that $\|\cdot\|$ is continuous in the metric $d(\cdot, \cdot)$. Having defined $\|\cdot\|$, we now state and prove the lemma.

LEMMA 4.1. *The norm $\|\cdot\|$ given by (1.8) is continuous in the metric $d(\cdot, \cdot)$ given by (2.2).*

PROOF. It suffices to show each of the norms $\|\cdot\|_L$, $\|\cdot\|_R$ and $\|\cdot\|_A$ is continuous in $d(\cdot, \cdot)$. The argument in each case is elementary. Noting that $d(x, x') < 1$ implies $z_n = z'_n$ for all n , and that $w_{n,i}^+ = w_{n,i} + \varepsilon_2$ is continuous in $w_{n,i}$, for given

i , the continuity of $\|\cdot\|_L$ is not difficult to see. For the same reasons and since $\psi_W(\cdot)$ and $\psi_A(\cdot)$ are locally Lipschitz, $\|\cdot\|_R$ and $\|\cdot\|_A$ are also continuous. \square

Basic inequalities. In this subsection, we demonstrate Propositions 4.1 and 4.2, which are at the foundation of the analysis in Sections 5 and 6. The evolution of $X(t)$ between arrivals of jobs is deterministic. In Proposition 4.1, we provide upper bounds on the rate of change of $\|X(t)\|$ there by employing its components $\|X(t)\|_L$, $\|X(t)\|_R$ and $\|X(t)\|_A$, which exist almost everywhere since the underlying functions are locally Lipschitz except where jobs arrive or depart, with jumps being negative at departures. We set

$$(4.15) \quad \overset{\circ}{\mu}_n = \begin{cases} \mu_n & \text{for class independent networks,} \\ 1 & \text{for station independent networks,} \end{cases}$$

and note that

$$(4.16) \quad \overset{\circ}{\mu}_n = \overset{\circ}{m}_{k,A} \mu_j.$$

For Proposition 4.1, as elsewhere in the paper, we assume the network is either class or station independent.

PROPOSITION 4.1. *For every subcritical JSQ network,*

$$(4.17) \quad \begin{aligned} \|X(t)\|'_L + \|X(t)\|'_R &\leq \|X(t)\|'_L + \frac{1}{2} \|X(t)\|'_R \\ &\leq \sum_n \overset{\circ}{\mu}_n (\varepsilon_1 - \psi_Z(Z_n(t))) \end{aligned}$$

for almost all t . Moreover, for any subset \mathcal{K} of $\{1, \dots, K\}$,

$$(4.18) \quad \sum_{k \in \mathcal{K}} \|X(t)\|'_{A,k} \leq (1 + \varepsilon_1/2) \rho_i \sum_n \overset{\circ}{\mu}_n \psi_Z(Z_n(t))$$

for almost all t , for $i = 1, 2$. Consequently, for almost all t ,

$$(4.19) \quad \|X(t)\|' \leq (\varepsilon_1/2) \sum_n \overset{\circ}{\mu}_n (2 - \psi_Z(Z_n(t))).$$

The inequality (4.19) follows immediately from (4.17) and (4.18), with $\mathcal{K} = \{1, \dots, K\}$, together with (1.8), since $\varepsilon_1 = 1 - \rho_i$. The purpose of the term $\frac{1}{2} \|X(t)\|'_R$ in the middle quantity in (4.17) is to permit a sharper upper bound, for large $\|X(t)\|_R$, by including the contribution from $\psi_W(\cdot)$ in (4.2). This strengthening of (4.19) will be applied in the proof of case (c) of Proposition 5.1. Also, since $\psi'_A(y) \rightarrow \infty$ as $y \rightarrow \infty$, the inclusion in the first sum in (4.18) of terms corresponding to certain $k \notin \mathcal{K}$ can improve the bound on the right-hand side, which will allow us to strengthen (4.19) for large $\|X(t)\|_A$.

PROOF OF PROPOSITION 4.1. We need to demonstrate the first two displays. For (4.17), we note that

$$\begin{aligned} & \|X(t)\|'_L + \frac{1}{2}\|X(t)\|'_R \\ & \leq -\sum_n \dot{\mu}_n \psi_Z(Z_n(t)) \sum_{i=1}^{Z_n(t)} \mathbf{1}\{W_{n,i}^+(t) < L_2\} R_{n,i}(t) \\ & \quad - \frac{1}{2} \sum_n \dot{\mu}_n \sum_{i=1}^{Z_n(t)} \psi'_W(W_{n,i}(t)) \mathbf{1}\{W_{n,i}^+(t) \geq L_2\} R_{n,i}(t) \\ & \leq -\sum_n \dot{\mu}_n \psi_Z(Z_n(t)) \sum_{i=1}^{Z_n(t)} \mathbf{1}\{W_{n,i}^+(t) < L_2\} R_{n,i}(t) \\ & \quad - \sum_n \dot{\mu}_n \psi_Z(Z_n(t)) \sum_{i=1}^{Z_n(t)} \mathbf{1}\{W_{n,i}^+(t) \geq L_2\} R_{n,i}(t) \\ & = -\sum_n \dot{\mu}_n \psi_Z(Z_n(t)) \mathbf{1}\{Z_n(t) > 0\} \leq \sum_n \dot{\mu}_n (\varepsilon_1 - \psi_Z(Z_n(t))) \end{aligned}$$

holds almost everywhere. The first inequality follows from the definitions of $\|\cdot\|_L$, $\|\cdot\|_R$ and (4.16), since $Z_n(\cdot)$ is constant almost everywhere. For this, note that $W'_{n,i}(t) = -\mu_{j_{n,i}} R_{n,i}(t)$ almost everywhere. The second inequality follows from (4.10), with the last inequality using $\psi_Z(0) = \varepsilon_1$ and $\sum_i R_{n,i}(t) = 1$, when $Z_n(t) > 0$. This implies (4.17).

For (4.18), we apply the definition of $\|\cdot\|_{A,k}$ in (4.11) to obtain, for almost all t ,

$$\begin{aligned} \sum_{k \in \mathcal{K}} \|X(t)\|'_{A,k} & \leq (1 + \varepsilon_1/2) \sum_{k \in \mathcal{K}} \sum_A \sum_n \alpha_k p_{k,A} q_{k,A,n} \dot{m}_{k,A} \psi_Z(Z_n(t)) \\ & \leq (1 + \varepsilon_1/2) \rho_i \sum_n \dot{\mu}_n \psi_Z(Z_n(t)), \end{aligned}$$

where the first inequality follows from $\psi'_A(y) \geq -1$, and the second inequality follows from the corollary to Lemma 3.1. \square

We still need to specify the constant ε_3 that was employed in (4.7); for later reference, we also recall the constant ε_2 from equations (4.4) and (4.6):

$$(4.20) \quad \varepsilon_2 = M_1 \dot{m}^{\text{ratio}} \varepsilon_3 = (\varepsilon_1)^2/40.$$

Recall that M_1 is specified in (4.12) and (4.13), and \dot{m}^{ratio} is defined as in the equation before (1.16). In Proposition 4.2, we show that, for this choice of ε_2 and ε_3 , the expected increase in $\|\cdot\|$ is nonpositive at the time T of the first arrival of a job in the network. We note that, for $X(0) = x$ fixed, T is deterministic, as is the evolution of $X(\cdot)$ up through $T -$.

PROPOSITION 4.2. *For every JSQ network,*

$$(4.21) \quad E_x[\|X(T)\|] \leq \|X(T-)\| \quad \text{for all } x.$$

PROOF. We consider the contribution to $\|\cdot\|$ from $\|\cdot\|_L$, $\|\cdot\|_R$ and $\|\cdot\|_A$, assuming that a single arrival occurs from stream k at time T ; when arrivals simultaneously occur from other streams, the corresponding bounds can be applied sequentially.

For $\|\cdot\|_L$, one has

$$(4.22) \quad \begin{aligned} & E_x[\|X(T)\|_L] - \|X(T-)\|_L \\ &= \sum_A \sum_n p_{k,A} q_{k,A,n}^* \sum_{i=1}^{Z_n(T-)} \mathring{m}_{k_n,i,A_n,i} (W_{n,i}^+(T) \wedge L_2) (\psi_Z(Z_n(T-) + 1) \\ & \quad - \psi_Z(Z_n(T-))) \\ & \quad + \sum_A \sum_n p_{k,A} q_{k,A,n}^* \mathring{m}_{k,A} \psi_Z(Z_n(T-) + 1) \\ & \quad \quad \times \int_0^\infty ((\mu_j y + \varepsilon_2) \wedge L_2) F_j(dy) \end{aligned}$$

for each x , where, for given k and A , $q_{k,A,n}^*$ is the probability that the arriving job is assigned to queue n . As previously, $j = (k, A, n)$. Because of the JSQ rule, $q_{k,A,n}^*$ is concentrated on the shortest queues in A . Since by (4.7),

$$(4.23) \quad \psi_Z(Z_n(T-) + 1) - \psi_Z(Z_n(T-)) \leq \varepsilon_3/L_2(Z_n(T-) + 1) \leq \varepsilon_3/L_2,$$

the first term on the right-hand side of (4.22) is at most

$$(4.24) \quad \varepsilon_3 \mathring{m}^{\text{ratio}} \sum_A p_{k,A} \mathring{m}_{k,A}.$$

On the other hand, since

$$\int_0^\infty ((\mu_j y + \varepsilon_2) \wedge L_2) F_j(dy) \leq \mu_j \int_0^\infty y F_j(dy) + \varepsilon_2 = 1 + \varepsilon_2,$$

which does not depend on n , and since $\psi_Z(y)$ is increasing in y and $q_{k,A,n}$ is concentrated on A , the last term on the right-hand side of (4.22) is at most

$$(4.25) \quad (1 + \varepsilon_2) \sum_A \sum_n p_{k,A} q_{k,A,n} \mathring{m}_{k,A} \psi_Z(Z_n(T-) + 1).$$

In other words, removing the truncation by L_2 in the integral and replacing $q_{k,A,n}^*$ (which is concentrated on the shortest queues in A) by $q_{k,A,n}$ can only increase this term. Note that this is the only place in the first six sections of the paper where the JSQ property is employed. It follows from the bounds (4.24) and (4.25) that

$$(4.26) \quad \begin{aligned} & E_x[\|X(T)\|_L] - \|X(T-)\|_L \\ & \leq (1 + \varepsilon_2) \sum_A \sum_n p_{k,A} q_{k,A,n} \mathring{m}_{k,A} (\psi_Z(Z_n(T-) + 1) + \mathring{m}^{\text{ratio}} \varepsilon_3). \end{aligned}$$

For $\|\cdot\|_R$, it follows from (4.2) and (4.4) that

$$\begin{aligned}
 & E_x[\|X(T)\|_R] - \|X(T-)\|_R \\
 (4.27) \quad &= \sum_A \sum_n p_{k,A} q_{k,A,n}^* \dot{m}_{k,A} \int_0^\infty \psi_W(\mu_j y) F_j(dy) \\
 &\leq \varepsilon_2 \sum_A p_{k,A} \dot{m}_{k,A}.
 \end{aligned}$$

On the other hand, it follows from (4.11) that

$$\begin{aligned}
 & E_x[\|X(T)\|_A] - \|X(T-)\|_A \\
 (4.28) \quad &\leq (1 + \varepsilon_1/2)(\varepsilon_3/L_2)M_1 \sum_A p_{k,A} \dot{m}_{k,A} \\
 &\quad - (1 - \varepsilon_2)(1 + \varepsilon_1/2) \sum_A \sum_n p_{k,A} q_{k,A,n} \dot{m}_{k,A} \psi_Z(Z_n(T-) + 1).
 \end{aligned}$$

In the first term, the factors ε_3/L_2 and M_1 are due to (4.23) and (4.12), since $\psi_A(0) = M_1$, and in the second term, the factor $1 - \varepsilon_2$ is due to (4.14).

Combining (4.26)–(4.28), it follows that

$$\begin{aligned}
 & E_x[\|X(T)\|] - \|X(T-)\| \\
 &\leq (4M_1 \dot{m}^{\text{ratio}} \varepsilon_3 + \varepsilon_2) \sum_A p_{k,A} \dot{m}_{k,A} \\
 &\quad - [(1 - \varepsilon_2)(1 + \varepsilon_1/2) - (1 + \varepsilon_2)] \\
 &\quad \times \sum_A \sum_n p_{k,A} q_{k,A,n} \dot{m}_{k,A} \psi_Z(Z_n(T-) + 1).
 \end{aligned}$$

Since $\psi_Z(y) \geq \varepsilon_1$ for all y , it follows from (4.20) that this is at most

$$-(\varepsilon_1^2/8) \sum_A p_{k,A} \dot{m}_{k,A} \leq 0.$$

So

$$E_x[\|X(T)\|] \leq \|X(T-)\|,$$

as desired. \square

The following upper bound is a consequence of Propositions 4.1 and 4.2 and the strong Markov property. It will be applied in the proof of Theorem 5.1.

COROLLARY 4.1. *For every subcritical JSQ network,*

$$(4.29) \quad E_x[\|X(t)\|] - \|x\| \leq C_2 t \quad \text{for all } t \text{ and } x,$$

where $C_2 = \sum_n \dot{\mu}_n$.

PROOF. Denoting by T_1, T_2, \dots the times at which arrivals occur and applying the strong Markov property, one can repeatedly apply Propositions 4.1 and 4.2 over the intervals $(0, T_1 \wedge t], (T_1 \wedge t, T_2 \wedge t], \dots$. Over each such interval, it follows from (4.19) and (4.21) that

$$(4.30) \quad \begin{aligned} & E_x[\|X(T_{i+1} \wedge t)\|] - E_x[\|X(T_i \wedge t)\|] \\ & \leq \left(\sum_n \overset{\circ}{\mu}_n\right) E_x[(T_{i+1} \wedge t) - (T_i \wedge t)] \end{aligned}$$

for each i , since $\psi_Z(y) > 0$ for all y . Summing over i gives

$$E_x[\|X(t)\|] - \|x\| \leq t \sum_n \overset{\circ}{\mu}_n,$$

and hence (4.29). \square

5. Proof of Theorem 1.1. In this section, we demonstrate Theorem 1.1. The proof is organized as follows. We first show it suffices to demonstrate Theorem 5.1, which is a slight variant of Theorem 1.1. The demonstration of Theorem 5.1 is then reduced to showing Proposition 5.1. The inequality in the proposition is expressed in terms of expected values of the norm $\|\cdot\|$ at stopping times σ that will be introduced shortly. Most of the rest of the section is devoted to showing Proposition 5.1, for which Propositions 4.1 and 4.2 are employed. At the end of the section, we briefly discuss a simpler alternative approach that can be applied to certain service disciplines; we also mention analogs of Theorem 1.1 where the state space is modified. We then discuss the stability of the JLLQ rule.

In order to demonstrate Theorem 1.1, we need to verify the inequality (1.10). In Theorem 5.1, we will instead demonstrate the variant (5.1). We recall that

$$\tau_M = \inf\{t : \|X(t)\| \leq M\}.$$

THEOREM 5.1. *For each subcritical queueing network satisfying (1.7), there exists M so that*

$$(5.1) \quad E_x[\tau_M] \leq C_3 \|x\| \quad \text{for all } x,$$

where $\|x\|$ is the norm given in (1.8) and $C_3 = (\varepsilon_1 \sum_n \overset{\circ}{\mu}_n)^{-1}$.

The inequality (1.10) follows quickly from Theorem 5.1 and Corollary 4.1. By (4.29),

$$E_x[\|X(1)\|] \leq \|x\| + C_2 \quad \text{for all } x,$$

where $C_2 = \sum_n \overset{\circ}{\mu}_n$. Restarting the process at time 1 and applying (5.1) to $x' = X(1)$ implies that

$$E_x[\tau_M(1)] \leq C_3(\|x\| + C_2) + 1,$$

which implies (1.10) with $C_1 = C_3 \vee (C_2 C_3 + 1)$.

The proof of Corollary 4.1 did not require any conditions on the evolution of $X(\cdot)$. In order to demonstrate Theorem 5.1, we need to consider the behavior of $X(\cdot)$ when its norm is large. If $Z_n(\cdot)$ is uniformly large over an interval for some n , we will be able to apply (4.19) of Proposition 4.1. If either $\|X(\cdot)\|_R$ or $\|X(\cdot)\|_A$ is large, we will be able to employ versions of (4.17) and (4.18). In each of these cases, we also apply Proposition 4.2. Iteration of these bounds and application of the strong Markov property as in the proof of the corollary will then imply the theorem.

In order to demonstrate (5.1), we need only consider $\|x\| > M$. On account of (1.8), $\|x\| > M$ implies that, for given M_L, M_R and M_A with

$$(5.2) \quad M = M_L + M_R + M_A,$$

either (a) $\|x\|_L > M_L$, (b) $\|x\|_L \leq M_L$ and $\|x\|_A > M_A$, or (c) $\|x\|_L \leq M_L$, $\|x\|_A \leq M_A$ and $\|x\|_R > M_R$. We will analyze these three cases for appropriate M_L, M_A and M_R .

Denote by T the time of the first arrival in the network. We introduce the stopping time σ , where

$$(5.3) \quad \sigma = \inf\{t : \|X(t)\|_L \leq M_L\} \wedge T$$

for x satisfying (a) and

$$(5.4) \quad \sigma = \inf\{t : \|X(t)\|_A \leq M_A\} \wedge T$$

for x satisfying (b). For x satisfying (c), we set

$$(5.5) \quad \sigma = t_x \wedge T'_{x, \Gamma+1}.$$

Here, t_x is deterministic and will be defined in (5.23); Γ is given in (1.7). (As mentioned there, $\Gamma = 0$ for many applications.) The term $T'_{x,i}$ is the time of the i th arrival at the queue n_x , with n_x being specified just before (5.23). We also set

$$\overset{\circ}{\mu}^{\text{ratio}} = \begin{cases} \max_n \overset{\circ}{\mu}_n / \min_n \overset{\circ}{\mu}_n & \text{for class independent networks,} \\ 1 & \text{for station independent networks,} \end{cases}$$

and recall L_1 , which was introduced in (4.8).

We will show the following proposition.

PROPOSITION 5.1. *For each subcritical JSQ network satisfying (1.7) and for L_1 satisfying $L_1 \geq 4N \overset{\circ}{\mu}^{\text{ratio}}$, there exist M_L, M_R and M_A , so that for x satisfying either (a), (b) or (c), and σ chosen as in (5.3)–(5.5),*

$$(5.6) \quad E_x[\sigma] \leq C_3(\|x\| - E_x[\|X(\sigma -)\|]),$$

where C_3 is as in Theorem 5.1.

Recall that $\|X(\cdot)\|$ has negative jumps at departures. It therefore follows from Proposition 4.2 and the strong Markov property that

$$(5.7) \quad E_x[\|X(\sigma)\|] \leq E_x[\|X(\sigma-)\|] \quad \text{for all } x.$$

Together with (5.6), this implies

$$(5.8) \quad E_x[\sigma] \leq C_3(\|x\| - E_x[\|X(\sigma)\|])$$

for x chosen as in the proposition.

It is not difficult to demonstrate Theorem 5.1 by iterating (5.8) and applying the strong Markov property.

PROOF OF THEOREM 5.1 USING (5.8). Iteration of (5.8), by applying the stopping rule σ at each step, induces a sequence of stopping times

$$0 < \sigma_1 < \sigma_2 < \dots,$$

with the sequence stopping at σ_I if $\|X(\sigma_I)\| \leq M$. Repeated application of (5.8), together with the strong Markov property, implies that, for each $i \leq I$,

$$(5.9) \quad E_x[\sigma_i] \leq C_3(\|x\| - E_x[\|X(\sigma_i)\|]) \leq C_3\|x\|$$

for all x . [The sum of the bounds obtained from the right-hand side of (5.8) forms a telescoping series.] On the other hand, over every finite time interval, there are only a finite number of arrivals and, in between arrivals, only the norm $\|\cdot\|_A$ can increase; hence only a finite number of stopping times can occur over a finite interval. It therefore follows from (5.9) and $\tau_M \leq \sigma_I$ that $\sigma_I < \infty$ almost surely, with

$$(5.10) \quad E_x[\tau_M] \leq E_x[\sigma_I] \leq C_3\|x\| \quad \text{for all } x.$$

This implies (5.1) of Theorem 5.1. \square

Most of the rest of this section is devoted to demonstrating Proposition 5.1. To do so, we consider separately the cases (a), (b) and (c) that are given after (5.2). Cases (a) and (b) will be easy to show; case (c) requires more effort. We employ the notation

$$(5.11) \quad \mathring{m}^{\max} = \max_{k,A} \mathring{m}_{k,A}.$$

For later use, we also set $\mathring{m}^{\min} = \min_{k,A} \mathring{m}_{k,A}$ and $\mathring{\mu}^{\min} = \min_n \mathring{\mu}_n$.

PROOF OF CASE (a) OF PROPOSITION 5.1. Under the condition (a), for each $t \in [0, \sigma)$, there exists an $n(t)$ so that $\|X(t)\|_{L,n(t)} > M_L/N$, and hence by the definition of $\|\cdot\|_{L,n}$ and by (4.9),

$$Z_{n(t)}(t) > M_L/(\mathring{m}^{\max} L_2 N \psi_Z(Z_{n(t)}(t))) \geq M_L/(\mathring{m}^{\max} L_1 L_2 N).$$

Setting

$$(5.12) \quad M_L = \overset{\circ}{m}^{\max} L_1 L_2 L_3 N,$$

it follows that $Z_{n(t)}(t) > L_3$, and hence $\psi_Z(Z_{n(t)}(t)) = L_1$. Since $L_1 \geq 4N\overset{\circ}{\mu}^{\text{ratio}}$, it follows from (4.19) of Proposition 4.1 that, almost everywhere on $(0, \sigma)$,

$$\|X(t)\|' \leq -(\varepsilon_1/2) \sum_n \overset{\circ}{\mu}_n (\psi_Z(Z_n(t)) - 2) \leq -\varepsilon_1 \sum_n \overset{\circ}{\mu}_n.$$

This implies case (a) of (5.6). \square

We next demonstrate case (b).

PROOF OF CASE (b) OF PROPOSITION 5.1. Under the condition (b), for each $t \in [0, \sigma)$, there exists a $k(t)$ so that $\|X(t)\|_{A,k(t)} > M_A/K$, and hence by the definition of $\|\cdot\|_{A,k}$ and by (4.9),

$$(5.13) \quad \psi_A(S_{k(t)}(t)) > M_A/(2\overset{\circ}{m}^{\max} L_1 K).$$

Choose $y_1 > M_1$ large enough so that

$$(5.14) \quad \psi'_A(y_1) \geq 2 \left(\sum_n \overset{\circ}{\mu}_n \right) / \min_{k,A} (\alpha_k \overset{\circ}{m}_{k,A})$$

and $\psi'_A(y_2) \geq \psi'_A(y_1)$ for $y_2 > y_1$; this is possible since $\psi'_A(y) \nearrow \infty$ as $y \nearrow \infty$.
Setting

$$(5.15) \quad M_A = 2\overset{\circ}{m}^{\max} L_1 K \psi_A(y_1),$$

it follows, from (5.13), that $\psi_A(S_{k(t)}(t)) > \psi_A(y_1)$ for each t , and hence that $S_{k(t)}(t) > y_1$ and $\psi'_A(S_{k(t)}(t)) \geq \psi'_A(y_1)$. For M_A as in (5.15), differentiation of $\|X(\cdot)\|_{A,k}$ using (4.11) and the lower bound ε_1 for $\psi_Z(y)$ therefore imply that, at $k = k(t)$ with $t \in [0, \sigma)$,

$$(5.16) \quad \|X(t)\|'_{A,k} \leq -\varepsilon_1 \sum_A \sum_n p_{k,A} q_{k,A,n} \alpha_k \overset{\circ}{m}_{k,A} \psi'_A(y_1),$$

which by (5.14) is at most $-2\varepsilon_1 \sum_n \overset{\circ}{\mu}_n$.

We apply (4.18) of Proposition 4.1, with \mathcal{K} equal to the complement of $\{k(t)\}$. Adding (4.17) to (4.18), one obtains the analog of (4.19), but with the additional term inherited from (5.16), namely, almost everywhere on $[0, \sigma)$,

$$(5.17) \quad \|X(t)\|' \leq -(\varepsilon_1/2) \sum_n \overset{\circ}{\mu}_n (2 + \psi_Z(Z_n(t))) \leq -\varepsilon_1 \sum_n \overset{\circ}{\mu}_n.$$

Case (b) of (5.6) follows. \square

We now begin the argument for case (c) of Proposition 5.1. This is the only part that requires the condition (1.7). It requires more work than the other two

cases since, when $\|x\|_R$ is large, we need sufficient service of some job (n, i) , with large $w_{n,i}$, to ensure a rapid decrease of $\|\cdot\|$. Since any service discipline is allowed, such service need not occur at all, or even most, times. We will instead show that for $\|x\|_L \leq M_L$, with M_L not too large, and $\|x\|_R > M_R$, with M_R enormous, there is a small (but, nevertheless, large enough) probability that a job (n, i) with enormous $w_{n,i}$ receives sufficient service so that the corresponding derivative $\psi'_W(w_{n,i})$ induces an enormous decrease over $(0, \sigma)$ of $\|\cdot\|_R$, and hence of $\|\cdot\|$. In particular, sufficient service of such a job (n, i) must occur if $(0, \sigma - m^{\max}]$ is sufficiently long to allow complete service of all jobs at n with smaller weighted residual service times. (Recall that $m^{\max} = \max_j m_j$.)

We will identify the job (n, i) , with “large $w_{n,i}$,” that was referred to in the last paragraph in terms of a rapidly increasing sequence $w(1), w(2), \dots$. This sequence will also be used to define t_x , which was used in the definition of σ in (5.5). We construct $w(i)$ and corresponding sequences $p(0), p(1), \dots$ and $t(0), t(1), \dots$ inductively.

Set

$$(5.18) \quad p(i) = h(t(i) + 1),$$

where $h(\cdot)$ is given in (1.7). We choose $w(i)$ and $t(i)$ so that

$$(5.19) \quad \psi'_W(w(i) - 1) = C_4 2^{i+\Gamma+2} (t(i - 1) + 1) / p(i - 1)$$

and

$$(5.20) \quad t(i) = \sum_{\ell=1}^i w(\ell) + 2\Gamma.$$

Here, $C_4 = \sum_n \hat{\mu}_n / \hat{\mu}^{\min}$, $\psi_W(\cdot)$ and Γ are as in (4.2) and (1.7), respectively. [Note that, on account of (4.4), the range of $\psi'_W(\cdot)$ contains $[1, \infty)$, and so (5.19) can always be solved for $w(i)$ and given i .] Often, $p(i)$ will decrease very rapidly to 0, and $w(i)$ and $t(i)$ will increase very rapidly. The factor 2^i in (5.19) is not required for the proof of Proposition 5.1, but will be used in the next section. Employing $w(i)$ and M_L , we set

$$(5.21) \quad M_R = \hat{m}^{\max} N \psi_W \left(\sum_{i=1}^{\lfloor C_5 M_L \rfloor} w(i) \right),$$

where $C_5 = 1/(\varepsilon_1 \varepsilon_2 \hat{m}^{\min})$.

The sequence $w(i)$ was defined in (5.19) with Lemma 5.1 in mind. For the lemma, we relabel the residual service times $w_{n,i}$, $i = 1, \dots, z_n$, in order of increasing value at each n , employing the notation $w'_{n,1}, w'_{n,2}, \dots, w'_{n,z_n}$. We employ $k'_{n,i}$ and $A'_{n,i}$ for the corresponding renewal streams and selection sets.

LEMMA 5.1. *Suppose that for given n , $\|x\|_{R,n} > M_R/N$ and $z_n \leq C_5 M_L$. Then, for some $i = 1, \dots, z_n$, $w'_{n,i} > w(i)$.*

PROOF. If $w'_{n,i} \leq w(i)$ for all $i = 1, \dots, z_n$, then

$$\begin{aligned}
 \|x\|_{R,n} &= \sum_{i=1}^{z_n} \mathring{m}_{k'_{n,i}, A'_{n,i}} \psi_W(w'_{n,i}) \leq \mathring{m}^{\max} \sum_{i=1}^{\lfloor C_5 M_L \rfloor} \psi_W(w(i)) \\
 (5.22) \qquad &\leq \mathring{m}^{\max} \psi_W\left(\sum_{i=1}^{\lfloor C_5 M_L \rfloor} w(i)\right) < \|x\|_{R,n},
 \end{aligned}$$

which is a contradiction. The first inequality holds since $z_n \leq C_5 M_L$, the second inequality since $\psi_W(\cdot)$ is convex with $\psi_W(0) = 0$, and the last inequality since $\|x\|_{R,n} > M_R/N$. \square

Suppose that for a given $x \in \mathcal{S}$, $\|x\|_{R,n} > M_R/N$ and $z_n \leq C_5 M_L$ for some queue n , which we denote by n_x . (In case of more than one such n , choose one of them.) As before Lemma 5.1, denote by $w'_{n_x,i}, i = 1, \dots, z_{n_x}$, the ordered sequence obtained from $w_{n_x,i}, i = 1, \dots, z_{n_x}$. We set

$$(5.23) \qquad t_x = m^{\max}(t(i_x - 1) + 1),$$

where i_x is the smallest index i at which $w'_{n_x,i} > w(i)$. On account of Lemma 5.1, such an index exists. The time t_x is used to define σ in (5.5) in case (c).

Using the preceding construction, we now complete the proof of Proposition 5.1.

PROOF OF CASE (c) OF PROPOSITION 5.1. Under case (c), $\|x\|_R > M_R$, and hence $\|x\|_{R,n} > M_R/N$ for some queue n . Moreover, since $\|x\|_L \leq M_L$,

$$(5.24) \qquad z_n \leq \|x\|_{L,n}/(\varepsilon_1 \varepsilon_2 \mathring{m}^{\min}) \leq M_L/(\varepsilon_1 \varepsilon_2 \mathring{m}^{\min}) = C_5 M_L,$$

with (4.6) and (4.7) being used for the first inequality. So, the assumptions of Lemma 5.1 are satisfied, and n_x, t_x and i_x can be defined as immediately following the lemma.

Let B_x denote the event where at most Γ potential arrivals occur at n_x by time t_x and their service times are each at most $2m^{\max}$. One can check that

$$(5.25) \qquad P_x(B_x) \geq 2^{-\Gamma} h(t_x/m^{\max}) = 2^{-\Gamma} p(i_x - 1),$$

where $h(\cdot)$ is as in (1.7). For the inequality, note that the probability of a service time being at most $2m^{\max}$ is at least $1/2$, and that the probability this occurs for all service times for up to Γ jobs is at least $1/2^\Gamma$, after conditioning on knowing the arrival streams and selection sets of the jobs. The equality follows from (5.18) and (5.23).

On B_x , $\sigma = t_x$. Moreover, by time t_x , the total service devoted to jobs at n_x , with $i \geq i_x$, is at least m^{\max} , since the total time required to serve all initial jobs i , with $i < i_x$, and the at most Γ arriving jobs is at most

$$(5.26) \qquad \sum_{i=1}^{i_x-1} m_{j'_{n_x,i}} w'_{n_x,i} + 2\Gamma m^{\max} \leq m^{\max} \left(\sum_{i=1}^{i_x-1} w(i) + 2\Gamma \right) = t_x - m^{\max}$$

for $j'_{n,i}$ defined analogously to $k'_{n,i}$ and $A'_{n,i}$, where the equality follows from (5.20) and (5.23). Since $\psi'_W(\cdot)$ is nondecreasing, it follows that, on B_x ,

$$(5.27) \quad \begin{aligned} \|x\|_{R,n_x} - \|X(t_x-)\|_{R,n_x} &\geq m^{\max} \overset{\circ}{\mu}^{\min} \psi'_W(w'_{n_x,i_x} - 1) \\ &\geq m^{\max} \overset{\circ}{\mu}^{\min} \psi'_W(w(i_x) - 1). \end{aligned}$$

Consequently,

$$(5.28) \quad \begin{aligned} E_x[\|x\|_{R,n_x} - \|X(\sigma-)\|_{R,n_x}; B_x] &= E_x[\|x\|_{R,n_x} - \|X(t_x-)\|_{R,n_x}; B_x] \\ &\geq m^{\max} \overset{\circ}{\mu}^{\min} 2^{-\Gamma} p(i_x - 1) \psi'_W(w(i_x) - 1) \\ &= C_4 \overset{\circ}{\mu}^{\min} 2^{i_x+2} t_x \geq 4C_4 \overset{\circ}{\mu}^{\min} t_x, \end{aligned}$$

with the first inequality following from (5.25) and (5.27), and the last equality following from (5.19).

On the other hand,

$$(5.29) \quad \begin{aligned} E_x[\|x\| - \|X(\sigma-)\|] - \frac{1}{2} E_x[\|x\|_{R,n_x} - \|X(\sigma-)\|_{R,n_x}; B_x] \\ \geq -\left(\sum_n \overset{\circ}{\mu}_n\right) E_x[\sigma]. \end{aligned}$$

To see this, note that, on intervals between arrivals, $\|X(t)\|_{R,n}$ is decreasing for all n , and so

$$(5.30) \quad \left[\|X(t)\| - \frac{1}{2}\|X(t)\|_{R,n_x} \mathbf{1}\{\omega \in B_x\}\right]' \leq \left[\|X(t)\| - \frac{1}{2}\|X(t)\|_{R}\right]'$$

almost everywhere. But at the time T_i of an arrival in the network,

$$(5.31) \quad \begin{aligned} &\left[\|X(T_i)\| - \frac{1}{2}\|X(T_i)\|_{R,n_x} \mathbf{1}\{\omega \in B_x\}\right] \\ &- \left[\|X(T_i-)\| - \frac{1}{2}\|X(T_i-)\|_{R,n_x} \mathbf{1}\{\omega \in B_x\}\right] \\ &\leq \|X(T_i)\| - \|X(T_i-)\| \end{aligned}$$

since $\|X(T_i)\|_{R,n_x} \geq \|X(T_i-)\|_{R,n_x}$.

One obtains

$$(5.32) \quad \left[\|X(t)\| - \frac{1}{2}\|X(t)\|_{R,n_x} \mathbf{1}\{\omega \in B_x\}\right]' \leq \sum_n \overset{\circ}{\mu}_n$$

by applying (4.17) and (4.18) of Proposition 4.1 to the right-hand side of (5.30) to get the analog of (4.19) [with $\frac{1}{2}\|X(t)\|'_R$, rather than $\|X(t)\|'_R$ being applied in (4.17)]. Also, by applying Proposition 4.2 to the conditional expectation with respect to \mathcal{F}_{T_i-} of the right-hand side of (5.31), one obtains

$$(5.33) \quad \begin{aligned} E_x[\|X(T_i)\| - \frac{1}{2}\|X(T_i)\|_{R,n_x} \mathbf{1}\{\omega \in B_x\}] \\ \leq E_x[\|X(T_i-)\| - \frac{1}{2}\|X(T_i-)\|_{R,n_x} \mathbf{1}\{\omega \in B_x\}]. \end{aligned}$$

One then obtains (5.29) from (5.32) and (5.33) and the strong Markov property by arguing as in the proof of Corollary 4.1.

It follows immediately from (5.28) and (5.29) that

$$\|x\| - E_x[\|X(\sigma-)\|] \geq 2C_4 \dot{\mu}^{\min} t_x - \left(\sum_n \dot{\mu}_n\right) E_x[\sigma].$$

Since $C_4 = (\sum_n \dot{\mu}_n) / \dot{\mu}^{\min}$ and $t_x \geq \sigma$, this implies

$$\|x\| - E_x[\|X(\sigma-)\|] \geq \left(\sum_n \dot{\mu}_n\right) E_x[\sigma],$$

which demonstrates case (c) of (5.6) of Proposition 5.1. \square

We note that, in some instances, the proof of case (c) of Proposition 5.1 is not needed, or can be simplified. For instance, if all of the service distributions have bounded support then, on account of (5.24), case (c) will be vacuous if M_R is chosen to be a large enough multiple of M_L , and so the proofs of the first two parts suffice.

Suppose, instead, that the service discipline is PS. Then, at any time t and queue n with $Z_n(t) \leq M_L$, $\|X(t)\|_{R,n}$ decreases at rate

$$(5.34) \quad \frac{\dot{\mu}_n}{Z_n(t)} \sum_{i=1}^{Z_n(t)} \psi'_W(W_{n,i}(t)) \geq \frac{\dot{\mu}_n}{M_L} \psi'_W(W_n(t)/M_L).$$

If M_L is fixed and M_R is chosen large enough then, for $\|X(t)\|_{R,n} > M_R$, the right-hand side of (5.34) is at least $4 \sum_{n'} \dot{\mu}_{n'}$ for some n . One can then argue analogously to the proof of case (b) that $\|X(t)\|$ decreases at least at rate $\sum_{n'} \dot{\mu}_{n'}$ for $\|X(t)\|_R > M_R$, which will imply (5.6) in case (c) as well. One can employ a similar argument for the FIFO service discipline, although one first needs to redefine the state space S so as to suppress information from the state space descriptor on the service times of jobs for which service has not yet begun. In none of these instances is the assumption (1.7) used, since lower bounds on the service effort devoted to jobs with large residual service times always hold. (Such lower bounds will not in general hold for LIFO and certain other standard service disciplines.)

When the interarrival times of a queueing network are exponentially distributed, one has the option of removing the residual interarrival times from the state space descriptor of the process $X(\cdot)$. The resulting process $X'(\cdot)$, which takes values x' in the corresponding space S' , will also be Markov under service disciplines that do not employ the omitted information. It is easy to see that positive Harris recurrence of $X(\cdot)$ implies the same for $X'(\cdot)$; its equilibrium is the projection of the equilibrium of $X(\cdot)$. If one wishes, one can instead demonstrate the analog of Theorem 1.1 directly for $X'(\cdot)$ by employing the norm $\|\cdot\|'$, with

$$(5.35) \quad \|x'\|' = \|x'\|_L + \|x'\|_R,$$

where $\|\cdot\|_L$ and $\|\cdot\|_R$ are defined as in (4.5) and (4.2). The proof simplifies somewhat, since one can combine Propositions 4.1 and 4.2, and one only requires the bounds (4.17), (4.26) and (4.27) given there. Since the state with no jobs is accessible and petite, positive Harris recurrence follows from this version of Theorem 1.1. One can also show the analog of Theorem 1.3, which is proved in the next section.

As mentioned in Section 2, one can enrich the space S by introducing other coordinates in the state space descriptor, such as (a) the elapsed time since the last arrival from each stream, or (b) the amount of service already received by each job. The arguments we have given for Theorems 1.1 and 1.2 remain the same in these settings, since the randomness of the networks—due to future interarrival and service times, and choices of the selection set—is not affected; nor is the norm $\|\cdot\|$. Theorem 1.3 can also be extended without difficulty to these settings, with the corresponding quantities being included in (1.24) and (1.25) if desired. It is easy to see from the summary of its proof that Corollary 1.1 holds under the enrichment (b); Corollary 1.1 also holds under (a), although an extra step is required at the end of the proof to match up the new coordinates. Similar comments hold for Corollary 1.2. In each of the above cases, when the interarrival times are exponentially distributed, one can simultaneously include new coordinates as in (b) while removing the coordinates corresponding to the residual interarrival time from the state space descriptor, as in the previous paragraph.

JLLQ networks. At the beginning of Section 1, we briefly mentioned JLLQ networks, where jobs are assigned to the queue with the smallest workload, that is, to the queue n where $\sum_{i=1}^n v_{n,i}$ is smallest. When two or more queues have the smallest workload, one of these queues is chosen according to some rule. The traffic intensities ρ_i , $i = 1, 2$, are defined as in (1.2) and (1.3) for the class and station independent cases, and the network is said to be subcritical if $\rho_i < 1$.

The stability of the network is not affected by the (nonidling) service discipline, since the evolution of the workload at a queue is not affected. Analogs of Theorems 1.1 and 1.2, and Corollary 1.1 for the stability of subcritical JLLQ networks hold, as do the uniform bounds in Theorem 1.3 and Corollary 1.2, with the assumptions (1.7) and (1.17) no longer being needed.

Stability of subcritical JLLQ networks is intuitively more obvious than for subcritical JSQ networks, since the system cannot be “tricked” into sending jobs into queues with high remaining work. Stability is also easier to show. In [11], it was shown by using the same argument involving fluid limits that was used for JSQ networks with the FIFO service discipline. One can also show stability, as well as uniform bounds on the marginal distributions, with the aid of an appropriate norm that satisfies the analog of Theorem 1.1. We do not supply a proof here, but provide some motivation.

Such a norm $\|\cdot\|$ is given by

$$(5.36) \quad \|x\| = \sum_{n=1}^N \psi(g_n),$$

with

$$(5.37) \quad g_n = \sum_{i=1}^{z_n} \mathring{m}_{k_n,i,A_n,i}^+ w_{n,i}^+ + (1 + \varepsilon_1/2) \sum_{k=1}^K \sum_{A \subseteq B_N} p_{k,A} q_{k,A,n} \mathring{m}_{k,A}^+ \psi_A(s_k).$$

The quantities $\mathring{m}_{k,A}^+$, w_n^+ , ε_1 , $p_{k,A}$, $q_{k,A,n}$, s_k and $\psi_A(\cdot)$ are the same as those that were employed in Section 4 to define the norm $\|\cdot\|$ there. The function $\psi : \mathbb{R}_{+,0} \rightarrow \mathbb{R}_{+,0}$ is required to be twice continuously differentiable, with $\psi(0) = 0$, $\psi'(y) > 0$, $\psi'(y) \rightarrow \infty$ as $y \rightarrow \infty$, $\psi''(y) > 0$, $\psi''(y) \rightarrow 0$ as $y \rightarrow \infty$, and

$$\int_0^\infty \psi(y) F_j(dy) < \infty \quad \text{for all } j.$$

The first sum in (5.37) plays a role similar to $\|\cdot\|_{L,n}$ for the JSQ rule, and the double sum in (5.37) plays a role similar to $\|\cdot\|_A$. (When the interarrival times are exponentially distributed, one can remove them from the state space descriptor, and omit the double sum.) The function $\psi(\cdot)$ has been chosen so that for large $\|X(t)\|$, $\|X(t)\|' < -C_6$ for given $C_6 > 0$. Since $\psi'(y) \rightarrow \infty$ as $y \rightarrow \infty$, the idling at empty queues does not affect this bound except in the computation of the constant. Also, since $\psi''(y) \rightarrow 0$ as $y \rightarrow \infty$, $\psi'(y)$ is ‘‘almost constant’’ locally for large y , which can be employed in conjunction with the subcriticality of the network to induce a negative drift for $\|X(t)\|$ at large values. We omit the details.

6. Uniform bounds on families of JSQ networks and tightness. In Theorem 1.2, we demonstrated positive Harris recurrence for the Markov process $X(\cdot)$ for subcritical JSQ networks, provided the sets A_M given there are petite. Such a network therefore has an equilibrium probability measure \mathcal{E} . Here, we consider the equilibria $\mathcal{E}^{(a)}$, $a \in \mathcal{A}$, of families \mathcal{A} of such networks, and demonstrate the uniform bounds on the tails of these equilibria that are given in (1.24) and (1.25) of Theorem 1.3. The proof of (1.24) occupies most of this section. After proving the theorem, we then apply it to show tightness and relative compactness of the marginal distributions under additional assumptions on the service disciplines.

We introduce the notation

$$(6.1) \quad \begin{aligned} \gamma_z^{(a)}(M) &= \mathcal{E}^{(a)}(Z_1 > M), & \gamma_\ell^{(a)}(M) &= \mathcal{E}^{(a)}(\mathcal{L}_1^{(a)} > M), \\ \gamma_w^{(a)}(M) &= \mathcal{E}^{(a)}(W_1^{(a)} > M) \end{aligned}$$

and

$$(6.2) \quad \gamma_s^{(a)}(M) = \max_k \mathcal{E}^{(a)}(S_k^{(a)} > M).$$

The quantities z_n , $\ell_n^{(a)}$, $w_n^{(a)}$ and $s_k^{(a)}$ were defined in (1.19)–(1.22); Z_n , $\mathcal{L}_n^{(a)}$, $W_n^{(a)}$ and $S_k^{(a)}$ are the corresponding random variables. (Recall that $\mathcal{L}_n^{(a)}$ is employed to avoid possible confusion with the constants L_i .) On account of the symmetry

condition (1.23), the probabilities in (6.1) do not depend on the specific queue n that is chosen; we denote by $z_1, \ell_1^{(a)}$ and $w_1^{(a)}$ the coordinates of a particular queue. To demonstrate (1.24) and (1.25) of Theorem 1.3, it suffices to show that

$$(6.3) \quad \gamma_z^{(a)}(M), \gamma_\ell^{(a)}(M), \gamma_w^{(a)}(M), \gamma_s^{(a)}(M) \rightarrow 0 \quad \text{as } M \rightarrow \infty,$$

uniformly in $a \in \mathcal{A}$.

The argument for $\gamma_s^{(a)}(M)$ is elementary. The arrival flows are renewal processes with densities $G_k^{(a)}(dy)$. The densities of the corresponding stationary distributions are therefore $\alpha_k y G_k^{(a)}(dy)$. Substituting this into (1.14) implies $\gamma_s^{(a)}(M) \rightarrow 0$ uniformly in a as $M \rightarrow \infty$, as desired.

The main idea in showing (6.3) for $\gamma_z^{(a)}(M)$ and $\gamma_w^{(a)}(M)$ will be to employ the bounds from Sections 4 and 5, for large $\|\cdot\|_{L,n}^{(a)}$ and $\|\cdot\|_{R,n}^{(a)}$, to show that if at a given queue n either the queue length z_n or the weighted workload $w_n^{(a)}$ is large with nonnegligible probability with respect to a given initial measure ν , then $E_\nu[\|X^{(a)}(t)\|^{(a)}]$ will decrease over an appropriate time interval. Since the measures $\mathcal{E}^{(a)}$ are stationary, this behavior will provide a contradiction unless (6.3) holds for both $\gamma_z^{(a)}(\cdot)$ and $\gamma_w^{(a)}(\cdot)$. We first demonstrate (6.3) for $\gamma_z^{(a)}(\cdot)$, which is not difficult. The argument for $\gamma_w^{(a)}(\cdot)$ is more involved, and relies on the argument for $\|\cdot\|_{R,n}$ in Section 5.

The limit in (6.3) for $\gamma_\ell^{(a)}(M)$ follows without difficulty from that for $\gamma_w^{(a)}(M)$. The basic idea is that jobs will not have time to age significantly before the queue empties, if the workload is typically low. We present this argument next.

PROOF OF (6.3) FOR $\gamma_\ell^{(a)}(\cdot)$. It suffices to show that, for each $\delta > 0$,

$$(6.4) \quad \gamma_\ell^{(a)}(M) \geq \delta \quad \Rightarrow \quad M \leq M_1(\delta)$$

for some function $M_1(\cdot)$ that does not depend on a . The argument involves partitioning the time interval $[0, g^{(a)}M]$, $g^{(a)} \stackrel{\text{def}}{=} (m^{\max})^{(a)}$, into subintervals of length $g^{(a)}M_2(\delta)$, for appropriate $M_2(\delta)$. One applies (6.3), for $\gamma_w^{(a)}(\cdot)$, and (1.17) to obtain lower bounds on the probability that a given queue is empty sometime on such an interval. Appropriate events corresponding to the intervals will be disjoint, and adding their probabilities will then imply (6.4).

We choose $M_2(\cdot)$ so that

$$(6.5) \quad \sup_{a \in \mathcal{A}} \gamma_w^{(a)}(M_2(\delta) - 2\Gamma) \leq \delta/2,$$

where Γ is as in (1.17). Also, denote by $A_i^{(a)}(\delta)$, $i = 0, 1, 2, \dots$, the events on which at most Γ potential arrivals occur at the queue over $(ig^{(a)}M_2(\delta), (i+1)g^{(a)}M_2(\delta))$, with each such arrival having service time at most $2g^{(a)}$. As in (5.25),

$$(6.6) \quad P_x(A_i^{(a)}(\delta)) \geq 2^{-\Gamma} h(M_2(\delta))$$

for all x and a . If we choose M so that $\gamma_\ell^{(a)}(M) \geq \delta$ for a given δ and a , it follows from (6.5), (6.6) and the stationarity of $\mathcal{E}^{(a)}$ that

$$(6.7) \quad \begin{aligned} P_{\mathcal{E}^{(a)}}(\mathcal{L}_1^{(a)}(iM_2(\delta)) > M, W_1^{(a)}(iM_2(\delta)) \leq M_2(\delta) - 2\Gamma; A_i^{(a)}(\delta)) \\ \geq \delta 2^{-\Gamma-1} h(M_2(\delta)). \end{aligned}$$

If at some time $\tau \in [ig^{(a)}M_2(\delta), g^{(a)}M]$, $Z_1^{(a)}(\tau) = 0$ holds then, for each $t \in [\tau, g^{(a)}M]$, it is immediate that $\mathcal{L}_1(t) \leq M$. On the other hand, under the event $B_i^{(a)}(\delta)$ on the left-hand side of (6.7), the total amount of work initially at or entering the queue over $[ig^{(a)}M_2(\delta), (i + 1)g^{(a)}M_2(\delta)]$ is at most $g^{(a)}(M_2(\delta) - 2\Gamma + 2\Gamma) = g^{(a)}M_2(\delta)$, which implies it will be empty at some time in $[ig^{(a)}M_2(\delta), (i + 1)g^{(a)}M_2(\delta)]$. Hence, on $B_i^{(a)}(\delta)$,

$$(6.8) \quad \mathcal{L}_1^{(a)}(t) \leq M \quad \text{for } t \in [(i + 1)g^{(a)}M_2(\delta), g^{(a)}M].$$

It follows from (6.8) and the definition of $B_i^{(a)}(\delta)$ that, for $\gamma_\ell^{(a)}(M) \geq \delta$ and $M \geq IM_2(\delta)$, $I \in \mathbb{Z}_+$, the events $B_i^{(a)}(\delta)$, $i = 0, \dots, I - 1$, are disjoint. Taking their union, it follows from (6.7) that

$$P_{\mathcal{E}^{(a)}}\left(\bigcup_{i=0}^{I-1} B_i^{(a)}(\delta)\right) \geq I\delta 2^{-\Gamma-1} h(M_2(\delta)).$$

Consequently, $I \leq 2^{\Gamma+1}/(\delta h(M_2(\delta)))$, and so, under $\gamma_\ell^{(a)}(M) \geq \delta$,

$$M \leq (2^{\Gamma+1}/(\delta h(M_2(\delta))) + 1)M_2(\delta) \stackrel{\text{def}}{=} M_1(\delta),$$

which does not depend on a . This implies (6.4). \square

Demonstration of (6.3) for $\gamma_z^{(a)}(\cdot)$. The uniformity conditions (1.13)–(1.16) ensure that the norms $\|\cdot\|^{(a)}$, $a \in \mathcal{A}$, can be defined, for a given L_1 , by choosing the quantities $\psi_Z(\cdot)$, $\psi_W(\cdot)$, $\psi_A(\cdot)$, L_2 , L_3 , M_1 , ε_1 , ε_2 and ε_3 from Section 4 so as not to depend on a . The uniform bounds obtained from these choices will be applied to show (6.3) for both $\gamma_z^{(a)}(\cdot)$ and $\gamma_w^{(a)}(\cdot)$. The term L_1 , $L_1 \geq 4$, should be thought of as a free variable; when showing (6.3), we will let $L_1 \rightarrow \infty$.

To see the above claim on the choice of these quantities, first note that there exist functions $\psi_W(\cdot)$ and $\psi_A(\cdot)$ satisfying the regularity and monotonicity properties given after (4.3) and (4.11), with $\psi_A(\cdot)$ satisfying (4.12) for appropriate M_1 , so that the following analogs of (4.4) and (4.13) hold:

$$(6.9) \quad \sup_{a \in \mathcal{A}} \max_j \int_0^\infty \psi_W(\mu_j^{(a)} y) F_j^{(a)}(dy) \leq \varepsilon_2$$

and

$$(6.10) \quad \sup_{a \in \mathcal{A}} \max_k \int_{M_1/\alpha_k^{(a)}}^\infty (\psi_A(\alpha_k^{(a)} y) + \alpha_k^{(a)} y) G_k^{(a)}(dy) \leq \varepsilon_2,$$

where $\varepsilon_2 = (\varepsilon_1)^2/40$ as in (4.20); ε_1 is determined by the left-hand side of (1.15). Both inequalities follow without difficulty from (1.13) and (1.14); the uniform limits of the tails in (1.13) and (1.14) permit the choice of $\psi_W(\cdot)$ and $\psi_A(\cdot)$ as in (6.9) and (6.10), with $\psi'_W(y) \nearrow \infty$ and $\psi'_A(y) \nearrow \infty$ as $y \nearrow \infty$.

As in (4.20), set $\varepsilon_3 = (\varepsilon_1)^2/(40M_1\hat{m}^{\text{ratio}})$, but with \hat{m}^{ratio} given by (1.16). For given L_1, L_2 and L_3 are chosen as in (4.8) and (4.9). The function $\psi_Z(\cdot)$ is then specified in (4.7), using these choices of $\varepsilon_1, \varepsilon_3, L_2$ and L_3 . Employing these quantities, one defines $\|\cdot\|_R^{(a)}, \|\cdot\|_L^{(a)}$ and $\|\cdot\|_A^{(a)}$ as in (4.2), (4.5) and (4.11), and $\|\cdot\|^{(a)}$ as in (1.18). These norms will depend on a in general.

Rather than deal directly with $\|\cdot\|^{(a)}$, we need to employ a truncated version in order to guarantee that its expectation with respect to $\mathcal{E}^{(a)}$ is finite. We denote by $\|\cdot\|^{(a,L_1)}$ the norm on $S^{(a)}$ for given L_1 and, by

$$(6.11) \quad \|x\|^{(a,\ell)} \stackrel{\text{def}}{=} \|x\|^{(a,L_1)} \wedge \bar{M} \quad \text{for } x \in S^{(a)},$$

its truncation at a given value \bar{M} , with $\ell \stackrel{\text{def}}{=} (L_1, \bar{M})$. (If the expectation of $\|\cdot\|^{(a,L_1)}$ with respect to $\mathcal{E}^{(a)}$, $a \in \mathcal{A}$, is finite, one can set $\bar{M} = \infty$.) Since $\mathcal{E}^{(a)}$, $a \in \mathcal{A}$, is invariant,

$$(6.12) \quad E_{\mathcal{E}^{(a)}}[\|X^{(a)}(t)\|^{(a,\ell)}] = E_{\mathcal{E}^{(a)}}[\|X^{(a)}(0)\|^{(a,\ell)}] \quad \text{for all } t.$$

As in Section 5, we denote by T the time of the first arrival in the network. Here and later on, when the context is clear, we drop the superscript (a) for quantities such as $X(\cdot)$ and $Z_n(\cdot)$ that are associated with the networks.

PROOF OF (6.3) FOR $\gamma_z^{(a)}(\cdot)$. We decompose

$$(6.13) \quad E_{\mathcal{E}^{(a)}}[\|X(0)\|^{(a,\ell)} - \|X(t)\|^{(a,\ell)}]$$

into two parts, depending on whether $T > t$ or $T \leq t$. For small t , the part with $T > t$ will contribute the main term. For $\|x\|^{(a,L_1)} \leq \bar{M}$ and given L_1 and \bar{M} , one has

$$(6.14) \quad \begin{aligned} & E_x[\|x\|^{(a,\ell)} - \|X(t)\|^{(a,\ell)}; T > t] \\ & \geq E_x[\|x\|^{(a,L_1)} - \|X(t)\|^{(a,L_1)}; T > t] \\ & \geq E_x\left[(\varepsilon_1\hat{\mu}^{(a)}/2) \sum_n \int_0^t (\psi_Z(Z_n(t')) - 2) dt'; T > t\right] \\ & \geq E_x\left[(\varepsilon_1\hat{\mu}^{(a)}/2) \sum_n \left((L_1/2) \int_0^t \mathbf{1}\{Z_n(t') > L_3\} dt' \right. \right. \\ & \quad \left. \left. - 2 \int_0^t \mathbf{1}\{Z_n(t') \leq L_3\} dt' \right); T > t\right], \end{aligned}$$

with the first inequality holding since $\|X(t)\|^{(a,\ell)} \leq \|X(t)\|^{(a,L_1)}$, the second inequality following from (4.19) of Proposition 4.1, and $\psi_Z(L_3) = L_1 \geq 4$ being

used for the third inequality. Here, $\overset{\circ}{\mu}^{(a)} \stackrel{\text{def}}{=} \overset{\circ}{\mu}_n^{(a)}$, which is constant in n by assumption.

On the other hand, for fixed a and L_1 ,

$$\mathcal{E}^{(a)}(\|X\|^{(a,L_1)} > \bar{M}) \rightarrow 0 \quad \text{as } \bar{M} \rightarrow \infty.$$

Also, $\|X(0)\|^{(a,\ell)} \geq \|X(t)\|^{(a,\ell)}$ for $\|X(0)\|^{(a,L_1)} > \bar{M}$, and the quantity inside $E_x[\cdot]$ for the last term in (6.14) is always at most $(\varepsilon_1 \overset{\circ}{\mu}^{(a)}/4)L_1 N^{(a)}t$. It therefore follows from (6.14) that

$$\begin{aligned} (6.15) \quad & E_{\mathcal{E}^{(a)}}[\|X(0)\|^{(a,\ell)} - \|X(t)\|^{(a,\ell)}; T > t] \\ & \geq E_{\mathcal{E}^{(a)}}\left[(\varepsilon_1 \overset{\circ}{\mu}^{(a)}/2) \sum_n \left((L_1/2) \int_0^t \mathbf{1}\{Z_n(t') > L_3\} dt' \right. \right. \\ & \quad \left. \left. - 2 \int_0^t \mathbf{1}\{Z_n(t') \leq L_3\} dt' \right); T > t \right] \\ & \quad - L_1 N^{(a)} t \delta_1^{(a)}(\bar{M}) \end{aligned}$$

for appropriate $\delta_1^{(a)}(\cdot)$, with $\delta_1^{(a)}(\bar{M}) \searrow 0$ as $\bar{M} \nearrow \infty$.

We now set $M = L_3$, for \bar{M} in (6.3). Since $Z_n(\cdot)$ has right continuous sample paths and so cannot immediately fall below L_3 if $Z_n(0) > L_3$, the right-hand side of (6.15) is at least

$$N^{(a)}t(\varepsilon_1 \overset{\circ}{\mu}^{(a)}((L_1/4)\gamma_z^{(a)}(L_3) - 1) - L_1 \delta_1^{(a)}(\bar{M}) - L_1 \delta_2^{(a)}(t)$$

for appropriate $\delta_2^{(a)}(\cdot)$, with $\delta_2^{(a)}(t) \searrow 0$ as $t \searrow 0$. Hence, for large enough \bar{M} and small enough t ,

$$\begin{aligned} (6.16) \quad & E_{\mathcal{E}^{(a)}}[\|X(0)\|^{(a,\ell)} - \|X(t)\|^{(a,\ell)}; T > t] \\ & \geq \varepsilon_1 \overset{\circ}{\mu}^{(a)} N^{(a)}t((L_1/4)\gamma_z^{(a)}(L_3) - 2). \end{aligned}$$

On the other hand, by employing (4.19) of Propositions 4.1, 4.2 and the strong Markov property, it follows that

$$(6.17) \quad E_{\mathcal{E}^{(a)}}[\|X(0)\|^{(a,\ell)} - \|X(t)\|^{(a,\ell)}; T \leq t] \geq -\varepsilon_1 \overset{\circ}{\mu}^{(a)} N^{(a)}t P_{\mathcal{E}^{(a)}}(T \leq t).$$

One can see this by considering the above difference over $[0, T]$ and $[T, t]$, and, for the second part, arguing as in the proof of Corollary 4.1.

Combining (6.16) and (6.17), and choosing \bar{M} large enough and t small enough, one obtains

$$(6.18) \quad E_{\mathcal{E}^{(a)}}[\|X(0)\|^{(a,\ell)} - \|X(t)\|^{(a,\ell)}] \geq \varepsilon_1 \overset{\circ}{\mu}^{(a)} N^{(a)}t((L_1/4)\gamma_z^{(a)}(L_3) - 3).$$

By the invariance of $\mathcal{E}^{(a)}$, the left-hand side of (6.18) is zero, and so

$$(6.19) \quad \gamma_z^{(a)}(L_3) \leq 12/L_1.$$

The limit (6.3) follows by letting $L_1 \rightarrow \infty$. \square

Demonstration of (6.3) for $\gamma_w^{(a)}(\cdot)$. On account of the limiting behavior in (6.3) for $\gamma_z^{(a)}(\cdot)$, it suffices to instead show that

$$(6.20) \quad \tilde{\gamma}_w^{(a)}(M) \rightarrow 0 \quad \text{as } M \rightarrow \infty$$

uniformly in a for

$$(6.21) \quad \tilde{\gamma}_w^{(a)}(M) \stackrel{\text{def}}{=} \mathcal{E}^{(a)}(W_1^{(a)} > M', Z_1 \leq M),$$

where M' is a function of M with $M' \rightarrow \infty$ as $M \rightarrow \infty$. We employ the same basic framework here as we did in analyzing $\gamma_z^{(a)}(\cdot)$, and will apply the truncated norm in (6.11) to (6.12), which we will show is violated unless the limit in (6.3) holds. As in the analysis of $\gamma_z^{(a)}(\cdot)$, we set $M = L_3$ here; we will set $M' = L_4$, which will be a function of L_3 (and hence of L_1) and will be defined in (6.23). As in the demonstration of (6.3) for $\gamma_z^{(a)}(\cdot)$, we continue to drop the superscript (a) when convenient and the context is clear.

The argument relies heavily on the constructions employed for the demonstration of case (c) of Proposition 5.1. We briefly recall the quantities that were employed and whether they depend on $a \in \mathcal{A}$. The quantities $\psi_Z(\cdot)$, $\psi_W(\cdot)$, $\psi_A(\cdot)$, L_2 , L_3 , M_1 , ε_1 , ε_2 and ε_3 were specified in the previous subsection and do not depend on a . As before, the quantity L_1 is allowed to vary, with L_2 and L_3 being functions of L_1 . We recall the terms Γ and $h(\cdot)$ from (1.17), and the terms $p(\cdot)$, $w(\cdot)$ and $t(\cdot)$ from (5.18)–(5.20), none of which depends on a . Instead of employing $C_4 = \sum_n \dot{\mu}_n / \dot{\mu}^{\min}$ as in (5.19), we now set

$$(6.22) \quad C_4 = L_3.$$

Also, rather than employing the quantities M_L and M_R from (5.12) and (5.21), we use L_3 and

$$(6.23) \quad L_4 \stackrel{\text{def}}{=} \psi_W \left(\sum_{i=1}^{L_3} w(i) \right).$$

Defining $w'_{n,i}$ as before Lemma 5.1, the conclusions of the lemma continue to hold if $w_n^{(a)} > L_4$ and $z_n \leq L_3$ replace the conditions on $\|x\|_{R,n}$ and z_n there, with the argument in the proof being the same. If $w_n^{(a)} > L_4$ and $z_n \leq L_3$, for given x and n , we define $t_{x,n}$ in terms of $i_{x,n}$ as in (5.23), but with $i_{x,n}$ being the smallest index i at which $w'_{n,i} > w(i)$.

Both $i_{x,n}$ and $t_{x,n}$ depend on x , and we wish to employ a fixed time that does not. For this, we note it follows from (6.21) that, for each $a \in \mathcal{A}$, there is a value $i^{(a)}$ at which

$$(6.24) \quad \mathcal{E}^{(a)}(i_{X,1} = i^{(a)}, W_1^{(a)} > L_4, Z_1 \leq L_3) \geq 2^{-i^{(a)}} \tilde{\gamma}_w^{(a)}(L_3).$$

After having chosen $i^{(a)}$, we set

$$(6.25) \quad t^{(a)} = (m^{\max})^{(a)}(t(i^{(a)} - 1) + 1).$$

We also denote by $B_n^{(a)}$, $n = 1, \dots, N$, the events where at most Γ potential arrivals occur at queue n by time $t^{(a)}$, their service times are each at most $2(m^{\max})^{(a)}$, $i_{X(0),n} = i^{(a)}$, $W_n^{(a)}(0) > L_4$ and $Z_n(0) \leq L_3$. It follows from (5.25) and (6.24) that

$$(6.26) \quad P_{\mathcal{E}^{(a)}}(B_n^{(a)}) \geq 2^{-i^{(a)}-\Gamma} p(i^{(a)} - 1) \tilde{\gamma}_w^{(a)}(L_3) \quad \text{for all } n.$$

In the following proof, $i^{(a)}$, $t^{(a)}$ and $B_n^{(a)}$ will assume the roles played by i_x , t_x and B_x in the proof of case (c) of Proposition 5.1.

PROOF OF (6.3) FOR $\gamma_w^{(a)}(\cdot)$. We first note that, for given \bar{M} ,

$$(6.27) \quad \begin{aligned} & E_{\mathcal{E}^{(a)}}[\|X(0)\|^{(a,\ell)} - \|X(t^{(a)}-)\|^{(a,\ell)}] \\ & \geq E_{\mathcal{E}^{(a)}}[\|X(0)\|^{(a,L_1)} - \|X(t^{(a)}-)\|^{(a,L_1)}; \|X(0)\|^{(a,L_1)} \leq \bar{M}] \end{aligned}$$

since $\|X(t^{(a)}-)\|^{(a,\ell)} \leq \|X(t^{(a)}-)\|^{(a,L_1)}$. Setting

$$(6.28) \quad \begin{aligned} \Delta_1^{(a)} &= \frac{1}{2} \sum_n E_{\mathcal{E}^{(a)}}[\|X(0)\|_{R,n}^{(a,L_1)} - \|X(t^{(a)}-)\|_{R,n}^{(a,L_1)}; \\ & \quad B_n^{(a)}, \|X(0)\|^{(a,L_1)} \leq \bar{M}] \end{aligned}$$

and

$$(6.29) \quad \begin{aligned} \Delta_2^{(a)} &= E_{\mathcal{E}^{(a)}}[\|X(0)\|^{(a,L_1)} - \|X(t^{(a)}-)\|^{(a,L_1)}; \\ & \quad \|X(0)\|^{(a,L_1)} \leq \bar{M}] - \Delta_1^{(a)}, \end{aligned}$$

we rewrite the right-hand side of (6.27) as

$$(6.30) \quad \begin{aligned} & E_{\mathcal{E}^{(a)}}[\|X(0)\|^{(a,L_1)} - \|X(t^{(a)}-)\|^{(a,L_1)}; \|X(0)\|^{(a,L_1)} \leq \bar{M}] \\ & = \Delta_1^{(a)} + \Delta_2^{(a)}. \end{aligned}$$

We first analyze $\Delta_1^{(a)}$, which will be the main term. The same reasoning as in (5.27) shows that $\Delta_1^{(a)}$ is at least

$$\frac{1}{2}(m^{\max})^{(a)} \mu^{(a)} N^{(a)} P_{\mathcal{E}^{(a)}}(B_1^{(a)}, \|X(0)\|^{(a,L_1)} \leq \bar{M})(\psi'_W(w(i^{(a)}) - 1)).$$

This is at least

$$(6.31) \quad \begin{aligned} & \frac{1}{2}(m^{\max})^{(a)} \mu^{(a)} N^{(a)} (2^{-i^{(a)}-\Gamma} p(i^{(a)} - 1) \tilde{\gamma}_w^{(a)}(L_3) - \delta_3^{(a)}(\bar{M})) \\ & \times \psi'_W(w(i^{(a)}) - 1) \\ & \geq C_4 \mu^{(a)} N^{(a)} \tilde{\gamma}_w^{(a)}(L_3) t^{(a)} \end{aligned}$$

for $\tilde{\gamma}_w^{(a)}(L_3) > 0$. The first line follows from (6.26), with $\delta_3^{(a)}(\bar{M}) \searrow 0$ as $\bar{M} \nearrow \infty$, and the following inequality follows [like the last equality in (5.28)] from (5.19),

for large enough \bar{M} . Combining the above inequalities, one obtains, for large enough \bar{M} ,

$$(6.32) \quad \Delta_1^{(a)} \geq C_4 \dot{\mu}^{(a)} N^{(a)} \tilde{\gamma}_w^{(a)}(L_3) t^{(a)}.$$

This holds trivially for $\tilde{\gamma}_w^{(a)}(L_3) = 0$.

We claim, on the other hand, that

$$(6.33) \quad \Delta_2^{(a)} \geq -\dot{\mu}^{(a)} N^{(a)} t^{(a)}.$$

The argument for this is essentially the same as that given for (5.29) in the proof of case (c) of Proposition 5.1. On intervals between arrivals, $\|X(t)\|_{R,n}$ is decreasing for all n , and so, for all ω ,

$$(6.34) \quad \left[\|X(t)\|^{(a,L_1)} - \frac{1}{2} \sum_n \|X(t)\|_{R,n}^{(a,L_1)} \mathbf{1}\{\omega \in B_n^{(a)}\} \right]' \leq \left[\|X(t)\|^{(a,L_1)} - \frac{1}{2} \|X(t)\|_R^{(a,L_1)} \right]'$$

holds almost everywhere. Also, at the time T_i of an arrival in the network,

$$(6.35) \quad \left[\|X(T_i)\|^{(a,L_1)} - \frac{1}{2} \sum_n \|X(T_i)\|_{R,n}^{(a,L_1)} \mathbf{1}\{\omega \in B_n^{(a)}\} \right] - \left[\|X(T_i-)\|^{(a,L_1)} - \frac{1}{2} \sum_n \|X(T_i-)\|_{R,n}^{(a,L_1)} \mathbf{1}\{\omega \in B_n^{(a)}\} \right] \leq \|X(T_i)\|^{(a,L_1)} - \|X(T_i-)\|^{(a,L_1)}$$

since $\|X(T_i)\|_{R,n}^{(a,L_1)} \geq \|X(T_i-)\|_{R,n}^{(a,L_1)}$. Arguing as after (5.32), one obtains the upper bound $\dot{\mu}^{(a)} N^{(a)}$ for the right-hand side of (6.34), and 0 for the expectation, over $\|X(0)\|^{(a,L_1)} \leq \bar{M}$, of the right-hand side of (6.35). Application of these bounds, together with the strong Markov property, will then imply (6.33).

Combining the bounds in (6.32) and (6.33), it follows from (6.27) that

$$(6.36) \quad E_{\mathcal{E}^{(a)}} [\|X(0)\|^{(a,\ell)} - \|X(t^{(a)}-)\|^{(a,\ell)}] \geq C_4 \dot{\mu}^{(a)} N^{(a)} \tilde{\gamma}_w^{(a)}(L_3) t^{(a)} - \dot{\mu}^{(a)} N^{(a)} t^{(a)}$$

for large \bar{M} . Since we have set $C_4 = L_3$, this is at least

$$(L_3 \tilde{\gamma}_w^{(a)}(L_3) - 1) \dot{\mu}^{(a)} N^{(a)} t^{(a)}.$$

On the other hand, since $\mathcal{E}^{(a)}$ is invariant, the left-hand side of (6.36) equals 0. Therefore,

$$\tilde{\gamma}_w^{(a)}(L_3) \leq 1/L_3,$$

and hence $\tilde{\gamma}_w^{(a)}(L_3) \rightarrow 0$ uniformly in a as $L_3 \rightarrow \infty$. This implies (6.20), and hence (6.3) for $\gamma_w^{(a)}(\cdot)$, as desired. \square

At the end of Section 5, it was noted that the proof of Theorem 1.1 simplifies for certain service disciplines, such as PS and FIFO. In particular, the proof of case (c) of Proposition 5.1 can be replaced by a simpler argument, as was outlined at the end of that section. The same is true for the demonstration of (6.3) for $\gamma_w^{(a)}(\cdot)$. For disciplines such as PS and FIFO, one can give a simpler argument along the lines of (6.3) for $\gamma_z^{(a)}(\cdot)$, by investigating the decrease in the expected value of the corresponding norms over a small enough time interval $[0, t]$. The reasoning is similar to that provided at the end of Section 5. Analogs of the other observations at the end of Section 5 also hold in the uniform setting of Section 6 as well.

Tightness and relative compactness of families of networks. We are interested here in the behavior of the projections of the equilibria measures $\mathcal{E}^{(a)}$ and $\bar{\mathcal{E}}^{(a)}$ onto $(S')^{(a)}$ and $(\bar{S}')^{(a)}$, $a \in \mathcal{A}$, for families \mathcal{A} of JSQ networks, with $K^{(a)} \equiv K$, for $(S')^{(a)}$ and $(\bar{S}')^{(a)}$ defined as in Section 2. The networks will be assumed to satisfy the hypotheses of Theorem 1.3. The measures $\bar{\mathcal{E}}^{(a)}$, $a \in \mathcal{A}$, are the natural extensions of $\mathcal{E}^{(a)}$, $a \in \mathcal{A}$, from $S^{(a)}$ to $\bar{S}^{(a)}$; they are concentrated on $S^{(a)}$. The spaces $(S')^{(a)}$ and $(\bar{S}')^{(a)}$ are assigned a fixed value N' , $N' \leq N^{(a)}$, for all $a \in \mathcal{A}$ (corresponding to the first N' queues); they are identical for different a , and hence can be denoted by S' and \bar{S}' . Such S' and \bar{S}' , which are projections of each $S^{(a)}$ and $\bar{S}^{(a)}$, will be referred to as *common projections* of the family \mathcal{A} . We denote the corresponding projected measures by $(\mathcal{E}^{(a)})'$ and $(\bar{\mathcal{E}}^{(a)})'$.

As examples of such families of networks, one can think of JSQ networks indexed by N , the number of queues for the network, with the number of arrival streams K being fixed. For $N \geq N'$, for given N' , the networks will share the common projection S' obtained by retaining only the first N' queues. As $N \rightarrow \infty$, one can investigate the limiting behavior of the projections $(\mathcal{E}^{(N)})'$ of the equilibria $\mathcal{E}^{(N)}$ onto S' .

To examine the common projections of $\mathcal{E}^{(a)}$ and $\bar{\mathcal{E}}^{(a)}$, for $a \in \mathcal{A}$, we recall from Section 2 the compact sets E_M in S and the compact sets \bar{E}_M in \bar{S} . The corresponding sets for the projections S' of S and \bar{S}' of \bar{S} will be compact with respect to their respective metrics; we also denote these sets by E_M and \bar{E}_M , respectively.

Under the assumptions of Theorem 1.3, the uniform limits (1.24) and (1.25) hold. The projections $(\bar{\mathcal{E}}^{(a)})'$ onto \bar{S}' of the measures $\bar{\mathcal{E}}^{(a)}$ are therefore *tight* with respect to the induced metrics. That is, for each $\varepsilon > 0$, there is a compact set $B \subset \bar{S}'$ so that $(\bar{\mathcal{E}}^{(a)})'(B) \geq 1 - \varepsilon$ for all $a \in \mathcal{A}$. The sets \bar{E}_M can be employed to see this: The conditions $z_n \leq M$, $\ell_{n,i}^{(a)} \leq M$ and $0 \leq w_{n,i}^{(a)} \leq M$ in the definition of \bar{E}_M clearly suffice for (1.24). On the other hand, the probability of at least one arrival in the network over $(0, 1/\alpha_k M)$ from a given arrival stream k is at most

$1/M$ for a network in equilibrium, and so the condition $1/M \leq s_k^{(a)} \leq M$ suffices for (1.25).

For certain service disciplines, the projections $(\mathcal{E}^{(a)})'$ onto S' of the measures $\mathcal{E}^{(a)}$ are also tight. In addition to employing the reasoning of the previous paragraph, one also needs to show that

$$(6.37) \quad \sup_{a \in \mathcal{A}} \mathcal{E}^{(a)}(W_{n,i}^{(a)} < 1/M \text{ for some } i, Z_n > 0) \rightarrow 0 \quad \text{as } M \rightarrow \infty$$

for given n . When the projections $(\mathcal{E}^{(a)})'$ are tight, it will be more informative to work with them than with the projections $(\bar{\mathcal{E}}^{(a)})'$, since all limits will be concentrated on S' .

Tightness of $(\mathcal{E}^{(a)})'$ is not difficult to show for PS networks since, when the number of jobs at a queue is bounded, each job must be served at at least a given rate. When there are jobs with scaled residual service times close to 0, these jobs will quickly leave the queue. On the other hand, the scaled rate at which jobs leave the queue is bounded in equilibrium, which therefore gives an upper bound on the expected number of jobs in equilibrium with scaled residual service times close to 0.

For FIFO networks, (6.37) will not be true in general since, depending on the choice of the network a , the distributions $F_j^{(a)}(\cdot)$ might be concentrated arbitrarily close to 0, and jobs that are not the oldest at their queue will not be served until the departure of older jobs. For families \mathcal{A} of networks where the service time distributions do not depend on a , (6.37) is not difficult to check, with the argument being similar to the argument for PS networks just mentioned. On the other hand, equation (6.37) will hold for general service distributions in the setting obtained by restricting the state space S by suppressing information on the service times of jobs for which service has not yet begun; this setting was mentioned at the end of Section 5. The argument in this setting proceeds as before.

Families of measures $(\mathcal{E}^{(a)})'$, $a \in \mathcal{A}$, need not be tight for arbitrary service disciplines. This is the case for LIFO service disciplines, even when $N^{(a)} = K^{(a)} \equiv 1$. For example, consider a family of networks $\mathcal{A} = \{5, 6, \dots, N, \dots\}$ having service distributions $F^{(N)}(\cdot)$ with $\mu^{(N)} \equiv 2$ and point masses of size at least $1/3$ at 1, and having interarrival distributions $G^{(N)}(\cdot)$ with $\alpha^{(N)} \equiv 1$ and support on $(0, 2]$, and with point masses of size at least $1/2$ at $1 - 1/N$. With probability at least $1/2$, an arriving job will immediately begin receiving service that continues until its residual service time is reduced to $1/N$, at which time its service is taken over by a new arrival. Using this, one can show that, for each N , the probability under the corresponding equilibrium measure $\mathcal{E}^{(N)}$ of there being at least one job with residual service time at most $1/N$ is at least $1/25$. This contradicts (6.37) and hence contradicts tightness.

If on the other hand, for a family of networks with LIFO service disciplines, the arrival streams are Poisson, it seems clear that (6.37) will hold, although the author does not see how to show this.

A family of probability measures on a metric space is *relatively compact* if, for each sequence drawn from the family, there is a subsequence that converges to some probability measure on the space. It follows from Prohorov’s theorem that a tight family of measures is automatically relatively compact (see, e.g., [2]). Combining this with the preceding discussion of the projected measures $(\bar{\mathcal{E}}^{(a)})'$ and $(\mathcal{E}^{(a)})'$ on common projections \bar{S}' , respectively S' , for a family of JSQ networks, one obtains the following conclusion from Theorem 1.3.

THEOREM 6.1. *Suppose that a family \mathcal{A} of JSQ networks, with $K^{(a)} \equiv K$, satisfies the uniformity conditions (1.13)–(1.17) and (1.23), and that for each network $a \in \mathcal{A}$, $A_M = \{x : \|x\|^{(a)} \leq M\}$ is petite for each $M > 0$ with respect to the norms in (1.18). Then the projected measures $(\bar{\mathcal{E}}^{(a)})'$ on each common projection \bar{S}' of the family \mathcal{A} are relatively compact. Moreover, if (6.37) also holds, then the projected measures $(\mathcal{E}^{(a)})'$ on each common projection S' are relatively compact.*

By restricting the family \mathcal{A} of networks under consideration, one obtains the following analog of Corollary 1.2.

COROLLARY 6.1. *Suppose that each member of a family \mathcal{A} of JSQ networks has a single Poisson arrival stream, that $F_j^{(a)}$ does not depend on j or a , that the selection rules are mean field and have uniformly bounded support, and that (1.26) holds. Then the projected measures $(\bar{\mathcal{E}}^{(a)})'$ on each common projection \bar{S}' of the family \mathcal{A} are relatively compact. Moreover, if (6.37) also holds, then the projected measures $(\mathcal{E}^{(a)})'$ on each common projection S' are relatively compact.*

7. A family of JSQ networks with large workload. In Section 6, we demonstrated Theorem 1.3. The limit (1.24) there states that in equilibrium, at each queue, the tails for the distribution on the number of jobs, their weighted ages and the weighted workload can be bounded uniformly for general families of networks. This bound does not depend on the service discipline. If one examines the proof of the theorem, one sees that the bounds that are obtained for the weighted workload are actually extremely weak. The term M' in the definition of $\tilde{\gamma}_w^{(a)}(\cdot)$ in (6.21) is given by $M' = L_4$, with L_4 being defined in (6.23) in terms of the sequence $w(1), w(2), \dots$ and L_3 . As noted after (5.21), $w(i)$ will often increase very rapidly. In fact, one can check that, for Poisson arrival streams and service distributions having any given number of moments, $w(i)$ can grow like

$$(7.1) \quad w(i + 1) \geq e^{bw(i)},$$

where $b > 0$ depends on the number of moments. The growth of M' in terms of M , with $M = L_3$, will therefore be far too rapid to infer anything useful quantitatively about the tail of the distribution of the weighted workload in equilibrium for a member of the family of networks.

Rather than providing quantitative information, the purpose of the uniform bounds in (1.24) of Theorem 1.3 was to establish tightness of the distributions on the number of jobs, weighted ages and weighted workload at a queue. This was discussed in Sections 1 and 6. One can, however, ask whether the rapid growth exhibited in (7.1), and hence by L_4 , is an artifact of the proof or whether similar bad behavior is in fact possible for the workloads of subcritical JSQ queueing networks. This point is of course relevant in deciding whether it is always advantageous to employ the JSQ rule for assigning arriving jobs, as opposed to, say, randomly choosing the queue, e.g., letting $D = 1$, in the setting of Section 1.

In this section, we present a family of networks for which the weighted workload exhibits bad behavior in an extreme manner that is of the order as that suggested by (7.1). The structure of the networks in the family is elementary in most aspects. Each network in the family possesses a single Poisson arrival stream and two queues to which jobs are directed; the state space S is therefore the same for each network. The selection set always consists of both queues, and the service distributions are both class and station independent. When the two queues have equal numbers of jobs, an arriving job will choose each of them with probability $\frac{1}{2}$. The rates of the Poisson arrivals will depend on the network as will the service distributions, but the traffic intensity for the different networks will be uniformly bounded away from one. The service discipline will be the same for all networks, but the discipline is concocted so as to produce inefficient service. In particular, service will be *nonlocal* in the sense that the choice of which jobs to serve will depend on the entire state of the network.

As we will see, this discipline will cause certain jobs with large service times to be served very slowly, with service being concentrated on the jobs with shorter service times at that queue, which therefore quickly leave the queue. The other queue will, during such times, serve its jobs with longer service times first, which causes it to accumulate jobs with short service times. The presence of these unserved jobs induces arriving jobs to be directed to the first queue. Since these arriving jobs at the first queue will be served before a job with large service time is served, this slows down the rate at which such a job is served. Before such a job with large service time receives all of its service, with high probability, a job with substantially greater service time will enter the queue, which causes the workload at the queue to grow. Iterating this behavior produces the growth in workload in which we are interested. This behavior occurs with few jobs at the queue in relation to its workload.

The remainder of the section is organized as follows. We first complete the description of the family of JSQ networks whose description was begun in the next to last paragraph. We then state Theorem 7.1, which describes the behavior of the weighted workload, in equilibrium, of these networks. The proof of the theorem consists of three parts and includes an induction argument on the size of a quantity related to the weighted workload.

The family of queues is indexed by ε , with $\varepsilon > 0$. Rather than directly give the service distributions and Poisson intensity of arrivals, we specify them as follows. The service distribution $F^{(\varepsilon)}(\cdot)$ of jobs at each network is assumed to be discrete, having point masses at $h(0), h(1), h(2), \dots$, with $h(0) = \delta \stackrel{\text{def}}{=} \gamma_0 \varepsilon$, for $\gamma_0 \in (0, \frac{1}{200}]$, $h(1) = 1$, $h(2) \in 2\mathbb{Z}_+$ with $h(2) \geq c_1$, where $c_1 \geq 100$ will be specified in the proof of Proposition 7.1 and does not depend on ε or γ_0 , $h(3) = (h(2))^3$, and

$$(7.2) \quad h(i + 1) = e^{\sqrt{h(i)}} \quad \text{for } i = 3, 4, \dots;$$

we will restrict the family to indices ε with $\varepsilon \leq 1/(h(3))^5$. [For convenience, we set $h(-1) = 0$.] We will refer to jobs with (initial) service time δ as *quick*, those with service time 1 as *moderate*, and those with service time $h(i)$, $i \geq 2$, as *large*. Jobs with these service times are assumed to arrive in the network at rate $2/\varepsilon$ for $h(0)$, $2(1 - \eta)$ for $h(1)$, where $\eta \in (0, \frac{1}{100}]$, and at rate $2(h(i))^{-i}$ for $h(2), h(3), \dots$.

One can check that the traffic intensity ρ is given by

$$(7.3) \quad \rho = \gamma_0 + 1 - \eta + \sum_{i=2}^{\infty} (h(i))^{1-i}.$$

We require that

$$(7.4) \quad \rho \leq 1 - \eta/2,$$

which is easy to do, because of the freedom in our choices for δ and $h(2)$. Note that the arriving moderate jobs require most of the work. Because of the growth of the exponents in the arrival rates for $h(i)$ as $i \rightarrow \infty$, the distribution functions $F^{(\varepsilon)}(\cdot)$ have all moments. Note that as $\varepsilon \searrow 0$, the mean of the service time goes to 0. This, together with the fixed arrival rates for moderate and large jobs, implies that the uniformity condition (1.13) in Theorem 1.3 is not satisfied.

We still need to specify the service discipline; as mentioned above, it depends on the entire state of the network. We first introduce some terminology. At each time, one of the queues will be the *designated queue*, with the other being the *other queue*. Provided the network is not empty, the designated queue will not be empty, in which case one of its jobs will be the *designated job*. We denote by $Y^{(\varepsilon)}(t)$ the residual service time at time t of the designated job; when both queues are empty, set $Y^{(\varepsilon)}(t) = 0$.

At $t = 0$, if the network is not empty, we choose one of the nonempty queues as the designated queue and a job with the largest residual service time for that queue as the designated job. This queue will remain the designated queue until it is empty and the other queue is not.

As time increases, a job becomes the designated job upon its arrival at the designated queue, if its service time is larger than or equal to the residual service time of the current designated job (and automatically becomes the designated job if the queue is empty). Only new jobs can become designated jobs, with the exception

being when service of a designated job is completed (that is, the job leaves the network). It will follow from the service discipline given in the next paragraph that the queue must then be empty; the other queue (if not empty) then becomes the designated queue, with the job with the largest residual service time at the queue becoming the designated job. In order to indicate the designated queue and job, an additional coordinate needs to be added to the state space S (which we continue to denote by S).

The allocation of service is specified as follows. At the designated queue, the designated job will only be served when there are no other jobs there. When a job arrives or departs from the network, among the nondesignated jobs, the job with the shortest residual service time at the designated queue receives all of the service at the queue; this service continues until the next arrival or departure from the network. At the other queue, upon an arrival or departure, the job with the longest residual service time receives all of the service, with the job continuing to receive this service until the next arrival or departure. When the designated queue and other queue switch, the service rules also switch. One can check that the designated job has the largest residual service time among jobs at its queue, although not necessarily among jobs in the entire network.

By (7.4), each member of the family of JSQ networks defined above is subcritical. Moreover, since the arrival stream is Poisson, the conditions (1.11) and (1.12) are satisfied. Consequently, by Corollary 1.1, the Markov process $X^{(\varepsilon)}(\cdot)$ underlying each network is positive Harris recurrent, and so has an equilibrium measure $\mathcal{E}^{(\varepsilon)}$. The following result gives a lower bound on the distribution of the sum $W_1^{(\varepsilon)} + W_2^{(\varepsilon)}$ of the weighted workloads of $\mathcal{E}^{(\varepsilon)}$. (As shown in the proof, the same bounds hold for the sum of the unweighted workloads.) Here, we set $\gamma_1 = \frac{1}{2000}$.

THEOREM 7.1. *For the family of networks defined above,*

$$(7.5) \quad \mathcal{E}^{(\varepsilon)}(W_1^{(\varepsilon)} + W_2^{(\varepsilon)} \leq h(\lfloor \gamma_1/\varepsilon \rfloor)) \leq 1/h(\lfloor \gamma_1/\varepsilon \rfloor).$$

On account of the recursion for $h(i)$ given by (7.2), $h(\lfloor \gamma_1/\varepsilon \rfloor)$ will be enormous when $1/\varepsilon$ is a moderate multiple of 2000, and in particular grows much more rapidly than $1/\varepsilon$; presumably, this rapid growth also holds for much smaller values than 2000. The square root in (7.2) is chosen for convenience; it can be replaced by any power strictly less than 1, if $h(3)$ is replaced by a correspondingly higher power of $h(2)$. Note that the rate of growth given here is somewhat slower than that of $w(i)$ in (7.1).

It follows from (7.5) that the weighted workload under $\mathcal{E}^{(\varepsilon)}$ is typically at least of order $h(\lfloor \gamma_1/\varepsilon \rfloor)$. This contrasts with the mean weighted workload for the equilibria of the networks obtained by setting $D = 1$, which is of order $1/\varepsilon$. (The mean workload for the equilibria of those networks is bounded over all ε .)

Theorem 7.1 holds because, when a designated queue has designated job with residual service time $h(i)$, a job with service time $h(i + 1)$ is more likely to arrive

at the queue before the residual service time of the designated job decreases to $h(i - 1)$, provided that $i \leq \gamma_1/\varepsilon$. One can therefore compare $X^{(\varepsilon)}(\cdot)$ with a discrete time birth–death process on $0, 1, \dots, \lfloor \gamma_1/\varepsilon \rfloor + 1$, with uniform positive drift. The comparison ceases to be valid when the designated queue has too large a multiple of $1/\varepsilon$ jobs.

Proof of Theorem 7.1. In order to make the preceding paragraph precise, we will employ a recursion argument that requires two propositions. The first proposition asserts that the above behavior occurs for $i = 2$; the second proposition employs the first proposition and an induction argument to show this behavior for general i . Since the proof of the first proposition is fairly long and involves a substantial number of estimates, we give a condensed proof of it.

For these propositions, we need to employ a modification of the process $X^{(\varepsilon)}(\cdot)$. We define a new family of Markov processes $X^{(\varepsilon, \kappa)}(\cdot)$, $\kappa = 0, 1, \dots, K^{(\varepsilon)}$, with $K^{(\varepsilon)} = \lfloor \gamma_1/\varepsilon \rfloor$, where $X^{(\varepsilon, \kappa)}(\cdot)$ evolves in the same manner as $X^{(\varepsilon)}(\cdot)$, but with the following modification of the JSQ rule: an arriving job at time t selects the queue with the smaller value of $Z_d^{(\varepsilon, \kappa)}(t-) + \kappa$ and $Z_o^{(\varepsilon, \kappa)}(t-)$, where $Z_d^{(\varepsilon, \kappa)}(\cdot)$ and $Z_o^{(\varepsilon, \kappa)}(\cdot)$ denote the number of jobs at the designated queue and other queue for the network indexed by (ε, κ) . The term κ will serve as a “handicap” for the designated queue. As before, we denote by $Y^{(\varepsilon, \kappa)}(\cdot)$ the residual service time of the designated job. Note that, for $\kappa = 0$, the queue selection rule is JSQ and $Y^{(\varepsilon, 0)}(\cdot) = Y^{(\varepsilon)}(\cdot)$.

Proposition 7.1 gives the following uniform bound on $Y^{(\varepsilon, \kappa)}(\cdot)$ over all $\kappa \leq K^{(\varepsilon)}$. We denote by y the residual service time of the designated job for a state $x \in S$.

PROPOSITION 7.1. *Suppose that $x \in S$ is any state for which $y = h(2)$. Let T denote the stopping time at which either $Y^{(\varepsilon, \kappa)}(T) = h(1)$ or $Y^{(\varepsilon, \kappa)}(T) = h(\ell)$, for $\ell \geq 3$, first occurs. Then, for $\kappa \leq K^{(\varepsilon)}$,*

$$(7.6) \quad P_x(Y^{(\varepsilon, \kappa)}(T) = h(1)) \leq 1/h(2).$$

The bound (7.6) depends strongly on our construction of $X^{(\varepsilon, \kappa)}(\cdot)$, where the designated queue tends to receive more jobs than the other queue.

Proposition 7.2 generalizes Proposition 7.1 to $h(i)$, $2 \leq i \leq K^{(\varepsilon)} + 1$, by applying Proposition 7.1 together with an induction argument. The proposition is stated for $\kappa \leq K^{(\varepsilon)} - i + 1$, although only the case $\kappa = 0$ is directly applied in the demonstration of Theorem 7.1. General κ are needed for the comparison in the first paragraph of the proof in the induction argument.

PROPOSITION 7.2. *Suppose that $x \in S$ is any state for which $y = h(i)$ for given i , with $2 \leq i \leq K^{(\varepsilon)} + 1$. Let T denote the stopping time at which either $Y^{(\varepsilon, \kappa)}(T) = h(i - 1)$ or $Y^{(\varepsilon, \kappa)}(T) = h(\ell)$, for $\ell \geq i + 1$, first occurs. Then, for $\kappa \leq K^{(\varepsilon)} - i + 1$,*

$$(7.7) \quad P_x(Y^{(\varepsilon, \kappa)}(T) = h(i - 1)) \leq 1/h(i).$$

We first prove Theorem 7.1, assuming Proposition 7.2. As mentioned earlier, the argument compares $Y^\varepsilon(\cdot)$ to a birth–death process.

PROOF OF THEOREM 7.1 ASSUMING PROPOSITION 7.2. Theorem 7.1 will follow from Proposition 7.2 by using the reasoning outlined in the paragraph before the beginning of this subsection. To see this, we first consider the sequence of stopping times $T(0), T(1), T(2), \dots$ that are defined inductively as follows, when $Y^{(\varepsilon)}(0) = h(i_0)$ for some $i_0 \geq 1$. We set $T_0 = 0$ and, when $i_0 \geq 2$, denote by $T(1)$ the first time at which either $Y^{(\varepsilon)}(T(1)) = h(i_0 - 1)$ or $Y^{(\varepsilon)}(T(1)) = h(\ell)$, for $\ell \geq i_0 + 1$. When $i_0 = 1$, $T(1)$ denotes the first time at which $Y^{(\varepsilon)}(T(1)) = h(\ell)$, for $\ell \geq 2$. The times $T(2), T(3), \dots$ are defined analogously.

For $n = 0, 1, 2, \dots$, we define the function $H_1^{(\varepsilon)}(n)$ by setting $H_1^{(\varepsilon)}(n) = i \wedge (K^{(\varepsilon)} + 1)$, where i is the value at which $Y^{(\varepsilon)}(T(n)) = h(i)$, and denote by $\mathcal{G}_n = \mathcal{F}_{T(n)}$ the σ -algebra at this time. It follows from the strong Markov property and Proposition 7.2, with $\kappa = 0$ and $2 \leq i \leq K^{(\varepsilon)} + 1$ that, for $H_1^{(\varepsilon)}(n) = i$,

$$(7.8) \quad P(H_1^{(\varepsilon)}(n + 1) = i - 1 \mid \mathcal{G}_n) \leq 1/4$$

for all ε ; when $i = 1$, the corresponding probability is 0.

One can compare $H_1^{(\varepsilon)}(\cdot)$ with the birth–death process $H_2^{(\varepsilon)}(\cdot)$ on $i = 1, 2, \dots, K^{(\varepsilon)} + 1$, where

$$P(H_2^{(\varepsilon)}(n + 1) = i - 1 \mid H_2^{(\varepsilon)}(n) = i) = 1/4,$$

$$P(H_2^{(\varepsilon)}(n + 1) = i + 1 \mid H_2^{(\varepsilon)}(n) = i) = 3/4$$

for $i \neq 1, K^{(\varepsilon)} + 1$, with the values i instead of $i - 1$ and i instead of $i + 1$ being taken at 1 and $K^{(\varepsilon)} + 1$, respectively. In particular, under $H_1^{(\varepsilon)}(0) = H_2^{(\varepsilon)}(0)$, one can couple $H_1^{(\varepsilon)}(\cdot)$ and $H_2^{(\varepsilon)}(\cdot)$ so that

$$(7.9) \quad H_1^{(\varepsilon)}(n) \geq H_2^{(\varepsilon)}(n) \quad \text{for all } n.$$

Note that $H_2^{(\varepsilon)}(\cdot)$ is Markov, but $H_1^{(\varepsilon)}(\cdot)$ is not.

Let $\mathcal{E}_2^{(\varepsilon)}$ denote the equilibrium measure on $\{1, 2, \dots, K^{(\varepsilon)} + 1\}$ of the process $H_2^{(\varepsilon)}(\cdot)$. The process is reversible with respect to $\mathcal{E}_2^{(\varepsilon)}$, with

$$(7.10) \quad \mathcal{E}_2^{(\varepsilon)}(B) = 1/3, \quad \mathcal{E}_2^{(\varepsilon)}(\{K^{(\varepsilon)} + 1\}) = 2/3,$$

where $B = \{1, 2, \dots, K^{(\varepsilon)}\}$.

Also, let $L_{N,1}^{(\varepsilon)}(A), L_{N,2}^{(\varepsilon)}(A)$ denote the number of visits for $H_1^{(\varepsilon)}(\cdot), H_2^{(\varepsilon)}(\cdot)$ to a set A , over $1, \dots, N$, starting from any initial state. By (7.9) and (7.10),

$$(7.11) \quad \limsup_{N \rightarrow \infty} \frac{1}{N} L_{N,1}^{(\varepsilon)}(B) \leq \lim_{N \rightarrow \infty} \frac{1}{N} L_{N,2}^{(\varepsilon)}(B) = 1/3,$$

$$\liminf_{N \rightarrow \infty} \frac{1}{N} L_{N,1}^{(\varepsilon)}(\{K^{(\varepsilon)} + 1\}) \geq \lim_{N \rightarrow \infty} \frac{1}{N} L_{N,2}^{(\varepsilon)}(\{K^{(\varepsilon)} + 1\}) = 2/3$$

both hold almost surely.

When \mathcal{G}_n is given, with $H_1^{(\varepsilon)}(n) = i \leq K^{(\varepsilon)}$, then $T(n + 1) - T(n)$ is stochastically dominated by an exponentially distributed random variable with mean $\frac{1}{2}h(i + 1)^{i+1}$. On the other hand, when \mathcal{G}_n is given, with $H_1^{(\varepsilon)}(n) = K^{(\varepsilon)} + 1$, then $T(n + 1) - T(n)$ stochastically dominates a random variable with mean $\frac{1}{8}h(K^{(\varepsilon)} + 2)^{K^{(\varepsilon)}+2}$. The upper bound is due to the rate at which jobs of size at least $h(i + 1)$ enter the network, as given below (7.2). The lower bound employs this rate, together with (7.8), to compare $T(n + 1) - T(n)$, restricted to the outcome $H_1^{(\varepsilon)}(n + 1) = K^{(\varepsilon)} + 1$, to the restriction of an exponentially distributed random variable.

We let

$$U_N^{(\varepsilon)}(A) = \int_0^{T(N)} \mathbf{1}\{Y^{(\varepsilon)}(t) \in A\} dt$$

denote the occupation time for $Y^{(\varepsilon)}(\cdot)$ of the set A up to the random time $T(N)$, starting from a given initial state. Recall that if $H_1^{(\varepsilon)}(n) = i$, then $Y^{(\varepsilon)} > h(i - 1)$ for $t \in [T(n), T(n + 1))$. It follows from this, (7.11), the two bounds in the previous paragraph, and the strong law applied to the incremental times $T(n + 1) - T(n)$ that

$$\begin{aligned} \limsup_{N \rightarrow \infty} \frac{1}{N} U_N^{(\varepsilon)}([0, h(K^{(\varepsilon)})]) &\leq \frac{1}{3} \cdot \frac{1}{2} \cdot h(K^{(\varepsilon)} + 1)^{K^{(\varepsilon)}+1}, \\ \liminf_{N \rightarrow \infty} \frac{1}{N} U_N^{(\varepsilon)}((h(K^{(\varepsilon)}), \infty)) &\geq \frac{2}{3} \cdot \frac{1}{8} \cdot h(K^{(\varepsilon)} + 2)^{K^{(\varepsilon)}+2} \end{aligned}$$

hold almost surely. Since $U_N^{(\varepsilon)}([0, \infty)) = T(N)$, it follows that

$$\begin{aligned} (7.12) \quad \limsup_{N \rightarrow \infty} \frac{U_N^{(\varepsilon)}([0, h(K^{(\varepsilon)})])}{T(N)} &\leq \frac{2h(K^{(\varepsilon)} + 1)^{K^{(\varepsilon)}+1}}{h(K^{(\varepsilon)} + 2)^{K^{(\varepsilon)}+2}} \\ &\leq 1/h(K^{(\varepsilon)} + 2) \leq 1/h(K^{(\varepsilon)}). \end{aligned}$$

Since the process $X^{(\varepsilon)}(\cdot)$ is ergodic, it follows from (7.12) that

$$(7.13) \quad \mathcal{E}^{(\varepsilon)}(Y^{(\varepsilon)} \leq h(K^{(\varepsilon)})) \leq 1/h(K^{(\varepsilon)}).$$

The residual service time $Y^{(\varepsilon)}$ of the designated job is bounded above by the (unweighted) workload in the network, which is at most $2\varepsilon(W_1^{(\varepsilon)} + W_2^{(\varepsilon)})$ since the mean service time is less than 2ε . Theorem 7.1 therefore follows from (7.13) and the definition of $K^{(\varepsilon)}$, after dropping the 2ε coefficient. \square

We now prove Proposition 7.1. The proof for the proposition is fairly long and requires a number of steps, where one obtains various bounds on the number of jobs entering each of the two queues. For these bounds, we will repeatedly employ

elementary large deviation bounds involving sums of independent random variables. We assume the reader is familiar with these types of estimates and sketch the corresponding steps. The first part of the proof, through (7.21), is given in detail.

CONDENSED PROOF OF PROPOSITION 7.1. Our basic reasoning will be as follows: When the other queue is not empty, there will typically be either moderate or large jobs there. On account of the service discipline at the queue, the moderate and large jobs will be served before the quick jobs, which allows quick jobs to accumulate. This contrasts with service at the designated queue, where quick jobs are served first and so do not accumulate. Therefore, because of the JSQ rule, when there are jobs in the other queue, most arriving jobs will select the designated queue. Since there will be jobs in the other queue most of the time, this implies that the amount of work in the designated queue will typically have a positive drift. The probability will therefore be small that $Y^{(\varepsilon, \kappa)}(t) \leq \frac{1}{2}h(2)$ before time $t_0 \stackrel{\text{def}}{=} (h(3))^4 = (h(2))^{12}$. But by time t_0 , there will, with high probability, be many arriving jobs in the network with service time at least $h(3)$, and hence at least one such large job will arrive at the designated queue. So (7.6) of the proposition will follow.

We first introduce some notation consisting of the events $A_1(n), A_2(n), \dots, A_8(n)$, with $n = 1, \dots, t_0$, and the stopping time T_0 . We denote by $A_1(n)$ the event on which the amount of work present at the designated queue at time n is at least $\frac{1}{2}h(2)$ and at all previous times is at least $\frac{1}{2}h(2) - 1$. Note that the designated queue remains the same over $[0, n]$ on $A_1(n)$. It is easy to see that, under $A_1(n)$, $Y^{(\varepsilon, \kappa)}(t) \geq \frac{1}{2}h(2) - 1$ for $t \leq n + 1$, that

$$(7.14) \quad A_1(n)^c = \emptyset \quad \text{for } n \leq \frac{1}{2}h(2),$$

and that $Y^{(\varepsilon, \kappa)}(T) = h(1)$, with $T \leq t_0$, can only occur when $A_1(t_0)^c$ occurs, where T is given in the statement of the proposition.

All of the other events $A_2(n), \dots, A_8(n)$ are assumed to be subsets of $A_1(n)$ that, in addition, satisfy the following properties. The event $A_2(n)$ occurs if the designated queue at no time in $[0, n]$ has more than $1/(1000\varepsilon)$ jobs; $A_3(n)$ is the event for which the other queue at no time in $[0, n]$ has more than $1/(500\varepsilon)$ jobs that were not present at time 0. The event $A_4(n)$ occurs if the subset of $[0, n]$ on which the other queue is empty has Lebesgue measure at most $\frac{5}{13}n$. The event $A_5(n)$ occurs if a total of at most $n/(200\varepsilon)$ jobs ever arrive at the other queue at times in $(0, n]$ at which there is a job in the queue with residual service time strictly greater than δ ; $A_6(n)$ is the event on which at most $\frac{1}{100}n$ of these arriving jobs are moderate or large. Let T_0 be the first time at which the other queue has at most $1/(500\varepsilon)$ jobs. The event $A_7(n)$ occurs if the subset of $[T_0, n]$ on which the other queue is not empty and has only jobs with residual service times at most δ has Lebesgue measure at most $\frac{1}{50}n$. The event $A_8(n)$ occurs if at least n moderate

or large jobs arrive at the designated queue over $(0, n]$. Except for (7.15), where $A_8(n)$ is employed, the precise definitions of the events $A_2(n), \dots, A_8(n)$ will not be needed until the paragraph after (7.21).

One can check that

$$(7.15) \quad A_1(n)^c \subseteq A_8(n-1)^c \quad \text{for } n \geq \frac{1}{2}h(2).$$

For this, note that $A_1(n-1)^c \subseteq A_8(n-1)^c$ and that if $\omega \in A_1(n-1) \cap A_1(n)^c$, then the amount of work present in the designated queue at time n is less than $\frac{1}{2}h(2)$, which implies $\omega \in A_8(n-1)^c$.

In order to show (7.6), we may assume that, at time 0, the designated queue has only a single job, consisting of a designated job with residual service time $h(2)$. Otherwise, one can wait until all jobs, except for the designated job, leave the queue, or a job with service time at least $h(3)$ arrives at the designated queue. Since the designated job is served last, one or the other event must eventually happen. If the former occurs, reset the time to 0.

Let $B(n)$ be the subset of $A_1(n)$ on which no job with service time at least $h(3)$ arrives at the designated queue by time n . Most of the work in showing (7.6) consists of showing the following bounds involving $B(t_0)$ and $A_\ell(n)$:

$$(7.16) \quad P_x(B(t_0)) \leq 2e^{-c_2h(2)}$$

and

$$(7.17) \quad P_x(A_1(n) \cap A_\ell(n)^c) \leq e^{-c_2h(2)} \quad \text{for } \ell = 2, \dots, 8$$

for integers $n \in [\frac{1}{2}h(2), t_0]$ and appropriate $c_2 > 0$ not depending on the choice of c_1 , which is given before (7.2). The inequalities (7.16) and (7.17), for $\ell = 8$, will be applied in the next paragraph. The other inequalities in (7.17) are needed to show this inequality. We assume that c_1 is chosen large enough so that $(c'_1)^{14} \leq e^{c'_1c_2}$ for $c'_1 \geq c_1$.

We now show (7.6) and afterwards motivate the inequalities. We may assume, by induction on n ,

$$(7.18) \quad P_x(A_1(n-1)^c) \leq (n-1)e^{-c_2h(2)};$$

(7.18) trivially holds for $n-1 \leq \frac{1}{2}h(2)$ on account of (7.14). Because of (7.15) and (7.17), with $\ell = 8$, it follows from (7.18) that

$$(7.19) \quad P_x(A_1(n)^c) \leq ne^{-c_2h(2)}.$$

Consequently, (7.19) holds for all $n \leq t_0$ and, in particular,

$$(7.20) \quad P_x(A_1(t_0)^c) \leq t_0e^{-c_2h(2)} = (h(2))^{12}e^{-c_2h(2)}.$$

As pointed out after (7.14), $Y^{(\varepsilon, \kappa)}(T) = h(1)$, with $T \leq t_0$, can only occur under this event.

We have so far not employed (7.16). By (7.16),

$$P_x(T > t_0; A_1(t_0)) \leq 2e^{-c_2h(2)}.$$

Together with (7.20), this implies that the event in (7.6) occurs with probability at most

$$(7.21) \quad 2(h(2))^{12}e^{-c_2h(2)} \leq 1/h(2),$$

with the inequality holding because of the assumption on c_1 given above and $h(2) \geq c_1$. This implies (7.6).

The rest of the proof of (7.6) is spent justifying (7.16) and (7.17). We begin by justifying (7.17), which requires most of the work; (7.16) will then follow quickly.

The exponential bounds in (7.17) follow from elementary large deviation bounds, involving sums of i.i.d. random variables, that can be obtained by using Markov’s inequality or by comparison with a birth–death process. We assume here that the reader is familiar with such bounds. Since the bounds for some of the previous indices ℓ are used to derive (7.17) for larger ℓ , the coefficient $c_{2,\ell}$ in the exponent of the inequality corresponding to (7.17) for a specific ℓ may depend on ℓ ; after computing this bound for each ℓ , one can then set $c_2 = \min_{\ell} c_{2,\ell}$ to obtain (7.17).

The inequality in (7.17) for $A_2(n)^c$ follows from elementary large deviations bounds and the JSQ rule (with handicap κ). To see this, note that there is initially only one job at the designated queue and (a) The number of moderate and large jobs entering the network will be of smaller order of magnitude than $1/(1000\varepsilon)$ off of a set of probability $e^{-h(2)}$ [since $1/(1000\varepsilon) \geq (h(3))^5/1000 \geq h(2)$ and $t_0 = (h(3))^4 \leq 10^{-6}(h(3))^5 \leq 10^{-6}/\varepsilon$]. (b) Work from quick jobs enters the network at rate at most $1/100$ (since $\gamma_0 \leq 1/200$), which is less than the rate 1 of service at the designated queue. Comparison with the corresponding birth–death process therefore implies that the net number of quick jobs entering and leaving the designated queue up to any time before t_0 is of smaller order than $1/(1000\varepsilon)$, off of a set of probability $e^{-h(2)}$ [since $t_0e^{-1/(1000\varepsilon)} \leq (h(3))^4 \exp\{-(h(3))^4\} \ll e^{-h(2)}$].

The inequality for $A_3(n)^c$ follows immediately from the inequality for $A_2(n)^c$, since it is assumed that the handicap κ is at most $K^\varepsilon \leq 1/(2000\varepsilon)$ and so $1/(1000\varepsilon) + K^\varepsilon + 1 \leq 1/(500\varepsilon)$.

In order to show the inequality (7.17) for $A_4(n)^c$, note that when the other queue is empty but the designated queue is not, the rate at which jobs with service time 1 arrive at the other queue is $2(1 - \eta) \geq 9/5$. When these jobs are being served, the other queue is not empty. Under $A_1(n)$, the designated queue is never empty on $[0, n]$. Noting that $1 - \frac{9}{5} \cdot \frac{5}{14} = \frac{5}{14}$, the frequency over $[0, n]$, for which the other queue is empty, is therefore typically at most $5/14$. The inequality in (7.17) follows from an elementary large deviations bound.

We next show (7.17) for $A_5(n)^c$. On $A_1(n)$, the other queue remains the same over $(0, n]$. Over this interval, the set of times having a job at the other queue, with

residual service time strictly greater than δ , is the union of disjoint intervals, each with length at least $2/3$, except for the first interval [since $h(1) = 1$ and $\delta < 1/3$]. The number of such intervals is at most $\frac{3}{2}n + 2 \leq 2n$. Note that no quick jobs are served at the other queue over such intervals. Also, at most n moderate or large jobs entering the other queue over $(0, n]$ can complete their service over $(0, n]$. It follows that, on $A_1(n) \cap A_3(n)$, there are at most $1/(500\varepsilon)$ jobs arriving at the other queue over any such interval, and hence at most $n/(250\varepsilon) + n \leq n/(200\varepsilon)$ jobs arriving at the other queue over such intervals. The inequality (7.17) for $A_5(n)^c$ therefore follows from that for $A_3(n)^c$.

Ordering these at most $n/(200\varepsilon)$ jobs in the order in which they arrive at the other queue, the service time of a job is independent of the service time of previous such jobs (although the total number of arriving jobs will not be independent). Since the probability a job is either moderate or large is less than ε , the probability of there being at least $n/100$ moderate and large jobs among $n/(200\varepsilon)$ jobs, $n \geq h(2)/2$, is exponentially small in $h(2)$. The inequality (7.17) for $A_6(n)^c$ therefore follows from that for $A_5(n)^c$.

For $A_7(n)^c$, note that there will typically be approximately $2n/\varepsilon$ arrivals in the network by time $n \geq \frac{1}{2}h(2)$ and so, by an elementary large deviations argument, there will be at least $3n/\varepsilon$ arrivals with an exponentially small probability in $h(2)$. At time T_0 , there are at most $1/(500\varepsilon)$ jobs at the other queue, and so there will be a total of at most $4n/\varepsilon$ jobs at the other queue after time T_0 . The amount of time spent serving a job after it has residual service time δ is of course δ , and so the time spent serving all of these jobs is at most $4(\delta/\varepsilon)n \leq \frac{1}{50}n$, which demonstrates (7.17) in this case as well.

We still need to show the inequality for $A_8(n)^c$. For this, note that by (7.17), with $\ell = 4$ and $\ell = 7$,

$$(7.22) \quad P_x(A_1(n) \cap (A_4(n)^c \cup A_7(n)^c)) \leq 2e^{-(c_{2,4} \wedge c_{2,7})h(2)},$$

and denote by $H \subseteq [T_0, n]$ the random set where there is at least one job at the other queue with residual service time strictly greater than δ . Then (7.22) gives an upper bound on the probability that $A_1(n)$ occurs and the measure of H is less than

$$n - T_0 - \left(\frac{5}{13}n + \frac{1}{50}n \right) \geq \frac{7}{12}n - T_0.$$

Setting $H' = (0, T_0] \cup H$, this implies $|H'| \geq \frac{7}{12}n$ off of the exceptional event. Since the service time of an arriving job is independent of the service times of previous jobs and $\eta \leq 1/100$, a large deviations bound implies that at least $\frac{8}{7}n$ moderate and large jobs enter the network on H' , off of an event of exponentially small probability in $h(2)$.

On the other hand, on $A_6(n)$, at most $\frac{1}{100}n$ moderate and large jobs arrive at the other queue during H . Also, on $A_2(n)$, no jobs can arrive at the other queue before time T_0 , since the handicap $\kappa \leq 1/(2000\varepsilon)$. So, on $A_2(n) \cap A_6(n)$, at most $\frac{1}{100}n$

moderate and large jobs arrive at the other queue during H' . Together with the last paragraph and (7.17), for $\ell = 2$ and $\ell = 6$, this implies that, off of an event of exponentially small probability in $h(2)$, when $A_1(n)$ occurs, at least $\frac{8}{7}n - \frac{1}{100}n \geq n$ medium and large jobs arrive at the designated queue over H' , and hence over $(0, n]$. This gives the inequality (7.17) for $A_8(n)^c$.

We now demonstrate (7.16). Setting $n = t_0$ and $\ell = 8$, we first note that, by (7.17),

$$(7.23) \quad P_x(A_1(t_0) \cap A_8(t_0)^c) \leq e^{-c_2 h(2)}.$$

On $A_8(t_0) \subseteq A_1(t_0)$, at least $t_0 = (h(3))^4$ medium and large jobs arrive at the designated queue over $(0, t_0]$. The probability such a job has service time at least $h(3)$ is greater than $(h(3))^{-3}$. So, when $(h(3))^4$ such jobs occur, the probability at least one of them has service time at least $h(3)$ is greater than $1 - 2e^{-h(3)}$. The bound (7.16) follows from this and (7.23). This completes the proof of the proposition. \square

We now prove Proposition 7.2. The argument involves repeated couplings.

PROOF OF PROPOSITION 7.2. We apply induction to demonstrate the proposition. The case where $i = 2$ was demonstrated in Proposition 7.1, so we still need to show that (7.7) holds for $i + 1$ given that it holds for i . The argument for this involves repeatedly coupling the process $X^{(\varepsilon, \kappa)}(\cdot)$, with $Y^{(\varepsilon, \kappa)}(0) = h(i + 1)$, to one of two processes $X^{(\varepsilon, \kappa)}(\cdot)$ and $X^{(\varepsilon, \kappa+1)}(\cdot)$, with $Y^{(\varepsilon, \kappa)}(0) = Y^{(\varepsilon, \kappa+1)}(0) = h(i)$. In our couplings, we refer to the networks corresponding to $X^{(\varepsilon, \kappa)}(\cdot)$, with $Y^{(\varepsilon, \kappa)}(0) = h(i + 1)$, on the one hand, and to the other two processes, on the other hand, as the *first* and *second* networks. As in the proof of Proposition 7.1, we may assume that, at time 0, the designated queue of the first network has only a single job, consisting of a designated job at $h(i + 1)$.

We compare the first network, with index (ε, κ) for given $\kappa \leq K^{(\varepsilon)} - i$, to the second network with the same index and having the exact same number of jobs with the exact same residual service times and designated job, except that the designated job of the second network has residual service time $h(i)$ instead of $h(i + 1)$. (Recall that the designated job is not required to have the largest residual service time in the network.) One can couple the processes so that their evolution is exactly the same, with respect to service of individual jobs and arrivals, until either (a) the residual service time of the designated job of the second network reaches $h(i - 1)$, (b) an arrival at the designated queue has service time $h(i)$, or (c) an arrival at the designated queue has service time at least $h(i + 1)$. In the last two cases, the arriving job becomes the designated job in the second network and, in the last case, it becomes the designated job in the first network.

By the induction assumption, the event in (a) occurs with probability at most $1/h(i)$. We consider such an event a *minor failure*. At the time σ at which it occurs, the designated job of the first network will have residual service time

$$h(i + 1) - h(i) + h(i - 1) \geq h(i + 1) - h(i).$$

When the event in (c) occurs, we refer to it as a *minor success*.

At the time σ the event in (b) occurs, we change the coupling. We now couple the first network to the second network with index $(\varepsilon, \kappa + 1)$, and having the same number of jobs with the same residual service times and designated queue as the first network, after ignoring the designated job of the first network. The second network therefore has one less job than the first network at time σ . The job that has arrived at time σ has service time $h(i)$ and so will be the designated job for the second network. It has the greatest residual service time for the designated queue in the first network, except for the designated job, and will be served last except for that job; it will be referred to as the *associate designated job* of the first network.

Since the first and second networks differ only by the presence of the large designated job in the former, and since the handicap of the second network is one greater than that for the first network, the networks can be coupled so that their evolution is the same until the time σ_1 at which either the event (a) or the event (c) in the second to last paragraph occurs. Under (a), a minor failure occurs whereas, under (c), a minor success occurs. The minor failure occurs with probability at most $1/h(i)$, at which time the associate designated job of the first network and the designated job of the second network each have residual service time $h(i - 1)$ and the designated job of the first network still has residual service time at least $h(i + 1) - h(i)$.

Combining the preceding two couplings, it follows from the above reasoning that the probability a minor failure occurs before a minor success is at most $2/h(i)$. One can repeat this reasoning up to $n(i)$ times, where

$$n(i) \stackrel{\text{def}}{=} \lfloor h(i + 1)/h(i) \rfloor - 2,$$

so that either (d) $n(i)$ minor failures or (e) a minor success has occurred. At this time σ_2 , the residual service time of the designated job of the first network is still greater than $h(i)$. The probability of the event in (d) is at most $(2/h(i))^{n(i)}$. If the residual service time of this job decreases to $h(i)$ before (e) occurs, we say a *failure* occurs.

When a minor success occurs, the arriving job has service time at least $h(i + 1)$ and becomes the designated job. If the service time of the arriving job is at least $h(i + 2)$, we call the minor success a *success* and stop the procedure. If a minor success that is not a success occurs, one can repeat the above comparisons, starting as before with a single job at the designated queue. Letting $Q(i)$ denote the number of minor successes occurring before a failure and B denote the event that a success occurs before a failure, it follows from the previous paragraph and the definition of $h(\cdot)$ that

$$\begin{aligned} P_x(Q(i) \leq e^{2\sqrt{h(i+1)}}, B^c) &\leq e^{2\sqrt{h(i+1)}} \left(\frac{2}{h(i)}\right)^{n(i)} \\ (7.24) \qquad \qquad \qquad &\leq \left(\frac{2e}{h(i)}\right)^{h(i+1)/h(i)} \leq \frac{1}{2h(i+1)}. \end{aligned}$$

On the other hand, the probability that a minor success will actually be a success is at least

$$(7.25) \quad \frac{1}{2}(h(i))^i / (h(i+1))^{i+1} \geq 1 / (h(i+1))^{i+1} \geq e^{-\sqrt{h(i+1)}}.$$

Together with (7.24), (7.25) implies that the probability of a failure occurring before a success is at most

$$\begin{aligned} & (1 - e^{-\sqrt{h(i+1)}})e^{2\sqrt{h(i+1)}} + \frac{1}{2h(i+1)} \\ & \leq \exp\{-e^{\sqrt{h(i+1)}}\} + \frac{1}{2h(i+1)} \leq \frac{1}{h(i+1)}. \end{aligned}$$

Since success and failure correspond to the two events in the statement of Proposition 7.2, this implies (7.7) of the proposition. \square

Acknowledgment. The author thanks an anonymous referee for a careful reading of the paper.

REFERENCES

- [1] AZAR, Y., BRODER, A., KARLIN, A. and UPFAL, E. (1994). Balanced allocations. In *Proc. 26th. ACM Symp. Theory Comp.* 593–602.
- [2] BILLINGSLEY, P. (1968). *Convergence of Probability Measures*. Wiley, New York. [MR0233396](#)
- [3] BRAMSON, M. (1998). Stability of two families of queueing networks and a discussion of fluid limits. *Queueing Systems Theory Appl.* **28** 7–31. [MR1628469](#)
- [4] BRAMSON, M. (2008). *Stability of Queueing Networks. Lecture Notes in Math.* **1950**. Springer, Berlin. [MR2445100](#)
- [5] BRAMSON, M. (2010). Network stability under max–min fair bandwidth sharing. *Ann. Appl. Probab.* **20** 1126–1176. [MR2680560](#)
- [6] BRAMSON, M., LU, Y. and PRABHAKAR, B. (2010). Randomized load balancing with general service time distributions. In *Proc. ACM SIGMETRICS 2010* 275–286. ACM, New York.
- [7] BRAMSON, M., LU, Y. and PRABHAKAR, B. (2011). Asymptotic independence of queues under randomized load balancing. *Queueing Systems Theory Appl.* Submitted.
- [8] DAI, J. G. (1995). On positive Harris recurrence of multiclass queueing networks: A unified approach via fluid limit models. *Ann. Appl. Probab.* **5** 49–77. [MR1325041](#)
- [9] DAVIS, M. H. A. (1993). *Markov Models and Optimization. Monographs on Statistics and Applied Probability* **49**. Chapman and Hall, London. [MR1283589](#)
- [10] FOLEY, R. D. and McDONALD, D. R. (2001). Join the shortest queue: Stability and exact asymptotics. *Ann. Appl. Probab.* **11** 569–607. [MR1865017](#)
- [11] FOSS, S. and CHERNOVA, N. (1998). On the stability of a partially accessible multi-station queue with state-dependent routing. *Queueing Systems Theory Appl.* **29** 55–73. [MR1643598](#)
- [12] LUCZAK, M. J. and MCDIARMID, C. (2005). On the power of two choices: Balls and bins in continuous time. *Ann. Appl. Probab.* **15** 1733–1764. [MR2152243](#)
- [13] LUCZAK, M. J. and MCDIARMID, C. (2006). On the maximum queue length in the supermarket model. *Ann. Probab.* **34** 493–527. [MR2223949](#)
- [14] MARTIN, J. B. and SUHOV, Y. M. (1999). Fast Jackson networks. *Ann. Appl. Probab.* **9** 854–870. [MR1722285](#)

- [15] MEYN, S. P. and TWEEDIE, R. L. (1993). Generalized resolvents and Harris recurrence of Markov processes. In *Doebelin and Modern Probability (Blaubeuren, 1991)*. *Contemp. Math.* **149** 227–250. Amer. Math. Soc., Providence, RI. [MR1229967](#)
- [16] MITZENMACHER, M. (2001). The power of two choices in randomized load balancing. *IEEE Transactions on Parallel and Distributed Systems* **12** 1094–1104.
- [17] NUMMELIN, E. (1984). *General Irreducible Markov Chains and Nonnegative Operators*. *Cambridge Tracts in Mathematics* **83**. Cambridge Univ. Press, Cambridge. [MR0776608](#)
- [18] OREY, S. (1971). *Lecture Notes on Limit Theorems for Markov Chain Transition Probabilities*. Van Nostrand Reinhold, London. [MR0324774](#)
- [19] SUHOV, Y. M. and VVEDENSKAYA, N. D. (2002). Fast Jackson networks with dynamic routing. *Problems Inf. Transmission* **38** 136–153. [MR2101275](#)
- [20] VÖCKING, B. (1999). How asymmetry helps load balancing. In *40th Annual Symposium on Foundations of Computer Science (New York, 1999)* 131–141. IEEE Computer Soc., Los Alamitos, CA. [MR1917553](#)
- [21] VVEDENSKAYA, N. D., DOBRUSHIN, R. L. and KARPELEVICH, F. I. (1996). Queueing system with selection of the shortest of two queues: An asymptotic approach. *Probl. Inform. Trans.* **32** 15–29. [MR1384927](#)
- [22] VVEDENSKAYA, N. D. and SUHOV, Y. M. (1997). Dobrushin’s mean-field approximation for a queue with dynamic routing. *Markov Process. Related Fields* **3** 493–526. [MR1607103](#)

TWIN CITIES CAMPUS
SCHOOL OF MATHEMATICS
INSTITUTE OF TECHNOLOGY
UNIVERSITY OF MINNESOTA
127 VINCENT HALL
206 CHURCH STREET SE
MINNEAPOLIS, MINNESOTA 55455
USA
E-MAIL: bramson@math.umn.edu