

DYNAMIC DISTRIBUTED SCHEDULING IN RANDOM ACCESS NETWORKS

ALEXANDER L. STOLYAR,* *Bell Laboratories*

Abstract

We consider a model of random access (slotted-aloha-type) communication networks of general topology. Assuming that network links receive exogenous arrivals of packets for transmission, we seek dynamic distributed random access strategies whose goal is to keep all network queues stable. We prove that two dynamic strategies, which we collectively call queue length based random access (QRA), ensure stability as long as the rates of exogenous arrival flows are within the network saturation rate region. The first strategy, QRA-I, can be viewed as a random-access-model counterpart of the max-weight scheduling rule, while the second strategy, QRA-II, is a counterpart of the exponential (EXP) rule. The two strategies induce different dynamics of the queues in the fluid scaling limit, which can be exploited for the quality-of-service control in applications.

Keywords: Random multiple access; queueing network; stability; medium access control; slotted aloha

2000 Mathematics Subject Classification: Primary 90B15; 60K25

Secondary 68M12

1. Introduction

In this paper we consider a model of random access (slotted-aloha-type) communication networks of general topology. Assuming that network links receive exogenous arrivals of packets for transmission, we seek dynamic distributed random access strategies whose goal is to keep the network queues stable. (Queues are formed by packets waiting for transmission at the links.) We prove that two dynamic strategies, which we collectively call queue length based random access (QRA), ensure stability as long as the rates of exogenous arrival flows are within the network *saturation rate region*.

Our basic network model (formally defined in Section 3) is that of [11], which in turn is a slight generalization of those in [12] and [21]. (A closely related, but different, model was considered in [10]. All the models mentioned above generalize the classical slotted aloha system [14].) Informally, our model consists of a finite set \mathcal{N} of nodes n and a finite set of directed communication links $i = (n, m)$ connecting some of the node pairs. Time is slotted. In each slot, node n transmits a ‘packet’ (or *accesses channel*) with probability p_n , and it chooses one of its outgoing links $i = (n, m)$ to transmit on, with probabilities p_{nm}/p_n summing up to 1. Link access probabilities p_i are, generally speaking, chosen dynamically, but once they are chosen for a slot, each node transmits independently of the others. A general interference structure between the nodes is defined, so that some of the simultaneous transmissions on the network links can fail, owing to the interference from other transmissions. Thus, a set

Received 19 October 2005; revision received 30 January 2008.

* Postal address: Bell Laboratories, Alcatel-Lucent, Murray Hill, NJ 07974, USA.

Email address: stolyar@research.bell-labs.com

of (constant) link access probabilities $p = \{p_i\}$ determines the corresponding set of link throughputs $\mu = \{\mu_i\}$ —this dependence is denoted by $\mu(p)$.

The network *saturation throughput region* M for the above model is the set of all vectors $\mu(p)$ for all possible access probability sets p , along with all nonnegative vectors x which can be dominated by such $\mu(p)$. In other words, region M is the set of all link throughput vectors which can be achieved by static random access strategies (i.e. those using constant access probabilities), under the assumption that all link queues are ‘saturated’ (i.e. have an unlimited number of packets available for transmission at any link at any time). Gupta and Stolyar [11] showed that the outer (Pareto) boundary of the saturation throughput region M was essentially characterized as the set of throughput vectors maximizing (over M) the system objective function $\sum_i w_i \log \mu_i$ for all possible sets $w = \{w_i\}$ of nonnegative weights assigned to the links. (See [10] for an analogous result for a related model.) An important property of the optimal solution μ^* for a given set of link weights w is that the corresponding set of access probabilities $p = p(w)$ (such that $\mu^* = \mu(p)$) can be determined in a distributed fashion; namely, to determine ‘its own’ access probabilities, each node will only need to ‘know’ the weights of the nodes in its appropriately defined ‘local neighborhood’. (This property was established earlier in [12] for the case when all $w_i = 1$.) It was also suggested in [11] (and [10]) that a network can be controlled by changing the link weights *dynamically* to satisfy some desired constraints of the link throughput allocation while keeping such allocation efficient (i.e. Pareto optimal). In fact, a specific algorithm was proposed in [11] which seeks to maximize the system utility, $\sum_i \alpha_i \log \mu_i$, subject to the desired lower bounds on the link rates, and was shown to have good performance.

In this paper we study the network model described above, but consider the case when each link i receives a flow of exogenous packet arrivals at the average rate λ_i . The question is: how can we set access probabilities dynamically so that all link queues are stochastically stable? Let $Q_i(t)$ be the queue length at link i in time slot t , and let the dynamic link weight $w_i(t)$ at time t be a function of $Q_i(t)$. A *QRA algorithm* uses (dynamic) link weights $w_i(t)$ and sets access probabilities in each slot t to $p(w(t))$. We consider two different variants of the QRA algorithm. The first algorithm, QRA-I, uses weights $w_i(t) = \alpha_i + \gamma_i Q_i(t)^\beta$ with parameters $\alpha_i \geq 0$, $\gamma_i > 0$, and $\beta > 0$. The second algorithm, QRA-II, uses weights $w_i(t) = \alpha_i \exp(\gamma_i Q_i(t)^\kappa)$, $\alpha_i > 0$, $\gamma_i > 0$, and $\kappa \in (0, 1)$. Our main result is that both the QRA-I and QRA-II algorithms ensure stability of the queues as long as the input rates vector $\lambda = \{\lambda_i\}$ of input rates lies (strictly) within the saturation throughput region M . We note that the QRA-I algorithm was essentially introduced in [11], but the algorithm stability issue was *not* addressed there; the QRA-II algorithm is new.

The reason why we consider two different algorithms is because, although they both ensure stability (as long as λ is within M), they induce different behavior of the queues. This is demonstrated by our analysis of their *fluid sample paths* (FSPs), which are roughly the limits $q(t)$, $t \geq 0$, of ‘fluid-scaled’ processes $(1/r)Q(rt)$, $t \geq 0$, as $r \rightarrow \infty$. We can say that QRA-I is the random-access-model counterpart of the much studied max-weight-type scheduling algorithms. (Max-weight algorithms were originally introduced in [20], and then extended and generalized to accommodate a large variety of models and scenarios (cf. [6] and [19] for recent reviews).) The dynamics of an FSP under QRA-I is such that a Lyapunov function of the form $\sum_i c_i \gamma_i q_i(t)^{\beta+1}$ is nonincreasing. In contrast, FSPs under the QRA-II algorithm are such that $\max_i \gamma_i q_i(t)$ is nonincreasing. In this sense, QRA-II is the counterpart of the ‘exponential’ (EXP) scheduling rule [16], which has the same property and is known to keep the values of $\gamma_i Q_i(t)$ roughly equal in the *heavy traffic* limit [17]. Thus, the QRA-II algorithm (unlike

QRA-I) allows us, to a certain extent, to ‘directly control’ the ratios of different queues—a useful feature for *quality-of-service* control in applications.

In terms of technique, the stability analysis of QRA-II is more involved than that of QRA-I. Similarly to the situation with the EXP algorithm, the conventional fluid scaling, leading to FSPs in the limit, is insufficient, and we additionally need to consider a ‘local fluid scaling’, leading to *local fluid sample paths* (LFSPs). Analysis of LFSP dynamics for QRA-II (in the proof of Theorem 4) is substantially different from previous analyses of LFSPs under the EXP algorithm in that it requires a completely different Lyapunov function. We believe that this part of our work is of independent interest.

We have to clarify why one would need QRA algorithms at all, and not simply use previously known max-weight and EXP algorithms for our model. Both max-weight and EXP algorithms would have the following form: in each time slot choose access probabilities so as to

$$\text{maximize } \sum_i w_i(t) \mu_i \quad (1)$$

with weights $w_i(t)$ defined as for QRA-I and QRA-II, respectively. Such an algorithm will ensure stability of the queues as long as the input rate vector λ lies within the system’s *maximum stability region* V . Region V is typically strictly larger than our saturation rate region M , because it is defined as the largest region within which stability is feasible at all under any strategy, including strategies utilizing global and instantaneous state sharing and coordination between network nodes. However, owing to the fact that ‘instantaneous service rates’ μ_i enter the sum in (1) *linearly*, solving (1) in each time slot would typically involve global (combinatorial) optimization, not allowing a distributed solution. In contrast, a QRA algorithm chooses access probabilities which

$$\text{maximize } \sum_i w_i(t) \log \mu_i,$$

and this *can* be done in a distributed fashion. Thus, although QRA algorithms guarantee stability within a typically smaller region of input rates, they are much easier to implement in practice.

We would like to mention one more line of previous research (see [1] and the references therein), which aims at characterizing the stability region S under static random access strategies (i.e. those with constant access probabilities) for the classical slotted aloha system (where all links interfere with each other). It is easy to see that the closure of S contains the saturation throughput region M , that is, $M \subseteq \bar{S}$. It has long been conjectured that, for slotted aloha, in fact $M = \bar{S}$. This conjecture was proved in [1], but only for the case when exogenous arrival processes are dependent in a special way.

The rest of the paper is organized as follows. In Section 2 we introduce the basic notation. Our network model is described in Section 3, and the queueing stability problem for the case of exogenous arrival processes is defined in Section 4. In Section 5 we define the saturation throughput region. The QRA algorithms and our main stability result (Theorem 1) are presented in Section 6. Section 7 contains the proof of Theorem 1; the key part of this proof (Theorem 2) is then proved separately for QRA-I and QRA-II in Sections 8 and 9, respectively.

2. Basic notation and conventions

We use the notation \mathbb{R} , \mathbb{R}_+ , and \mathbb{R}_{++} for the sets of real, real nonnegative, and real positive numbers, respectively. Corresponding I -times product spaces are denoted \mathbb{R}^I , \mathbb{R}_+^I , and \mathbb{R}_{++}^I . The space \mathbb{R}^I is viewed as a standard vector space, with elements $x \in \mathbb{R}^I$ being row vectors

$x = (x_1, \dots, x_I)$, $x \cdot y$ being the scalar product, and $\|x\| = (x \cdot x)^{1/2}$ being the Euclidean norm, inducing the standard metric.

Vector equalities and inequalities are understood componentwise. Sometimes, where it cannot cause confusion, we slightly abuse notation by applying \log and \exp componentwise,

$$\log x = (\log x_1, \dots, \log x_I), \quad \exp x = (\exp x_1, \dots, \exp x_I),$$

and by writing

$$\gamma \times q = (\gamma_1 q_1, \dots, \gamma_I q_I)$$

for the componentwise product of vectors.

For a scalar function $F(x)$, $x \in \mathbb{R}^I$, and a subset $M \in \mathbb{R}^I$, $x^* \in \arg \max_{x \in M} F(x)$ means that x^* maximizes $F(x)$ within M . The abbreviation ‘u.o.c.’ means ‘uniform on compact sets’ convergence of functions. Usually, we will consider this convergence for functions (or vector functions) of the time $t \in [0, \infty)$, in which case u.o.c. convergence means uniform convergence in $[0, b]$ for any $b \geq 0$. Define $[z]^- := \min\{z, 0\}$ and $[z]^+ := \max\{z, 0\}$, and, for a nonnegative a , define $[z]_a^+$ as z if $a > 0$ and $[z]_a^+$ as $[z]^+$ if $a = 0$.

3. Basic model

We consider the basic model of [11], which is a generalized version of the model of [12] and [21]. The system consists of a finite set $\mathcal{N} = \{1, 2, \dots, N\}$ of nodes and operates in discrete time, with time slots indexed by $t = 0, 1, 2, \dots$. Let $\mathcal{D}_n \subseteq \mathcal{N} \setminus n$ denote the subset of nodes to which node n has data to send. A node n at any time t may attempt transmission of one unit of data, called a packet or *customer*, to one of the nodes $m \in \mathcal{D}_n$. When this happens, we say that node n makes a transmission attempt on the *link* (n, m) . We will denote by

$$\mathcal{I} := \{(n, m) \mid n \in \mathcal{N}, m \in \mathcal{D}_n\}$$

the set of all system links, and by I its cardinality (i.e. the total number of links). Throughout the paper, for brevity, we often (but not always) denote links $(n, m) \in \mathcal{I}$ by a single index i .

We make the following additional model assumptions.

- (a1) A node cannot simultaneously (i.e. within the same slot) transmit on two or more different links.
- (a2) If a node transmits in a slot, any simultaneous attempt to transmit to this node will fail.
- (a3) If there are two or more simultaneous transmissions to a node, they all collide and fail.
- (a4) Any transmission attempt by node n will interfere with and ‘erase’ any attempt to receive a packet at any of the nodes within some subset of \mathcal{N} , denoted by \mathcal{N}_n . (The model of [12] and [21] additionally assumed that $m \in \mathcal{N}_n$ implied that $n \in \mathcal{N}_m$.)

Note that assumptions (a3) and (a4) imply that $\mathcal{D}_n \subseteq \mathcal{N}_n$. (In other words, a transmission attempt by node n may interfere with receiving at more nodes than it actually sends traffic to.) Also, by assumption (a2), $n \in \mathcal{N}_n$ for all n .

We consider a class of *random access* (slotted-aloha-type) transmission schemes, defined as follows. In each time slot t each node n chooses the set of *link access probabilities* p_{nm} , $m \in \mathcal{D}_n$, such that

$$p_n := \sum_{m \in \mathcal{D}_n} p_{nm} \leq 1. \tag{2}$$

Node n attempts a transmission of one packet (customer) in slot t with probability p_n (which can be called the node access probability) independently of other nodes, and when it does transmit, it chooses a particular link $m \in \mathcal{D}_n$ to transmit on, also randomly, with probabilities p_{nm}/p_n . The resulting link transmission success probabilities (or, average link *throughputs*) in the slot are given by

$$\mu_{nm} = p_{nm} \prod_{\{k: m \in \mathcal{N}_k, k \neq n\}} (1 - p_k). \tag{3}$$

We emphasize that the access probabilities p_{nm} may depend on time, past history, and be mutually dependent across nodes and links. However, once the p_{nm} s are chosen for a given slot, the node transmission attempts are independent.

If the random access scheme is *static*, that is, probabilities p_{nm} stay constant in time, the μ_{nm} s are time average link throughputs.

The dependence of the set (vector) of link throughputs $\mu = (\mu_i, i \in \mathcal{I}) \in \mathbb{R}_+^I$ on the set (vector) of link access probabilities

$$p \in \mathcal{P} := \{(p_i, i \in \mathcal{I}) \in [0, 1]^I \mid (2) \text{ holds}\},$$

given by (3), will be denoted by $\mu(p)$. Clearly, the function $\mu(p)$ is continuous.

4. Stability problem

Suppose that there is an exogenous arrival process of packets (customers) of average rate λ_i to be transmitted on link i . To simplify exposition, we will assume that the arrival processes for different links are mutually independent, and that the process for link i is given by an independent and identically distributed (i.i.d.) sequence $A_i(t)$, $t = 0, 1, 2, \dots$, of nonnegative integer random variables, where $A_i(t)$ is the number of packets (customers) arriving in slot t . (Obviously, $\lambda_i = E[A_i(t)]$.) Also, for simplicity, we assume that $P\{A_i(t) = 0\} > 0$. (These assumptions on the arrival processes can be replaced by much more general Markov assumptions, e.g. those in [2]. Essentially, all we will need is that the underlying stochastic process describing the evolution of the system, under the strategy we will introduce later in Section 6, is Markov.)

Customers waiting for transmission form queues: one queue each link. The queue length (i.e. number of waiting customers) at link i at time t is denoted by $Q_i(t)$.

The problem is to find a dynamic random access strategy, i.e. a dynamic rule for choosing access probabilities p_i , such that the link queues remain stochastically stable. For a random access strategy such that the queueing process $Q(t) = (Q_i(t), i \in \mathcal{I}), t \geq 0$, is a Markov chain (as will be the case for the QRA strategies we introduce later in Section 6); stochastic stability is understood as ergodicity of this Markov chain.

Obviously, stability cannot be expected for arbitrary input rates $\lambda = (\lambda_i, i \in \mathcal{I})$. The QRA strategies of Section 6 are such that they ensure stability, as long as λ lies within the saturation throughput region, which we introduce next.

5. Saturation throughput region

We define the system's saturation throughput region M , which we will often simply call the *throughput region*, as the set of all nonnegative vectors, which can be majorized by vectors of the form $\mu(p)$, namely,

$$M := \{\mu' \in [0, 1]^I \mid \mu' \leq \mu(p) \text{ for some } p \in \mathcal{P}\}.$$

Region M has the following simple interpretation. Suppose that each link is ‘saturated’, i.e. there is an unlimited number of packets available for transmission. A vector μ' is within M if there exists a static random access strategy, with some constant vector p of access probabilities, that provides average throughput of at least μ'_i on each link i . It is easy to observe that region M is a compact subset of the positive quadrant \mathbb{R}_+^I .

It was shown in [11] that the subset of maximal elements (or Pareto boundary) of M , i.e. the set

$$M^* := \{\mu^* \in M \mid \mu^* \leq \mu' \in M \text{ implies that } \mu' = \mu^*\},$$

is essentially equal to the set of points $\mu \in M$ maximizing (‘weighted proportional fairness’) objective $\sum_{i \in \mathcal{I}} w_i \log \mu_i$ for different nonnegative weights w_i . The choice of access probabilities p such that $\mu(p)$ maximizes $\sum w_i \log \mu_i$ over M is given by Proposition 1, below, which is Theorem 1 of [11], which in turn is a generalization of the corresponding result in [12].

For each $n \in \mathcal{N}$, let us denote by

$$\mathcal{S}_n := \{(\ell, k) \mid k \in \mathcal{D}_\ell, k \in \mathcal{N}_n\}$$

the set of all links (ℓ, k) which either originate at n or are such that a transmission by node n interferes with that on (ℓ, k) .

Proposition 1. ([11, Theorem 1].) *For an arbitrary set of positive weights*

$$w = (w_{nm}, (n, m) \in \mathcal{I}) \in \mathbb{R}_{++}^I,$$

there exists a unique set of access probabilities $p \in \mathcal{P}$ which maximizes the function

$$\sum_{(n,m) \in \mathcal{I}} w_{nm} \log \mu_{nm}(p).$$

The optimal p is given by

$$p_{nm} = \frac{w_{nm}}{\sum_{(\ell,k) \in \mathcal{S}_n} w_{\ell k}}. \tag{4}$$

The dependence of the set (vector) of access probabilities $p \in \mathcal{P}$ on the set (vector) of positive link weights $w \in \mathbb{R}_{++}^I$, given by (4), will be denoted by $p(w)$. We will extend the domain of $p(w)$ for all $w \in \mathbb{R}_+^I$, using the convention that $p_{nm} = 0$ when $w_{nm} = 0$. (Expression (4) is well defined when $w_{nm} > 0$, but may not be when $w_{nm} = 0$.) Clearly, $p(w)$ is continuous at any point $w \in \mathbb{R}_{++}^I$. It is not necessarily continuous at points w with some zero components. However, some continuity properties, namely those in Proposition 2(c), below, do hold and will suffice for our purposes.

Proposition 2. (a) *Function $p(w)$, $w \in \mathbb{R}_+^I$, is invariant with respect to the scaling of w by a positive constant.*

(b) *Suppose that $w' \in \mathbb{R}_+^I$. Then,*

$$\mu' \in \arg \max_{x \in M} \sum w'_{nm} \log x_{nm} \text{ if and only if } \mu'_{nm} = \mu_{nm}(p(w'))$$

for all (n, m) with $w'_{nm} > 0$. (We use the convention that $0(-\infty) = 0$.)

(c) Suppose that $w \rightarrow w' \in \mathbb{R}_+^I$. Then the following assertions hold.

(c1) If $w'_{nm} = 0$ and $w'_{\ell k} > 0$ for at least one $(\ell, k) \in \mathcal{S}_n$ then $p_{nm}(w) \rightarrow 0 = p_{nm}(w')$. In particular, if $w'_{nm} = 0$ and $w'_{nk} > 0$ then $p_{nm}(w) \rightarrow 0$.

(c2) If $w'_{nm} > 0$ then

$$p_{nm}(w) \rightarrow p_{nm}(w') \quad \text{and} \quad \mu_{nm}(p(w)) \rightarrow \mu_{nm}(p(w')).$$

The proof is straightforward and is thus omitted.

Remark. Expression (4) can be equivalently rewritten as

$$p_{nm} = \frac{w_{nm}}{\sum_{k \in \mathcal{N}_n} W_k^{in}}, \tag{5}$$

where

$$W_k^{in} := \sum_{\ell: k \in \mathcal{D}_\ell} w_{\ell k}$$

is the sum of the weights of all links ‘incoming’ to node k . As explained in [11], given the set of weights w , the calculation of access probabilities p_{nm} according to (5) can be done in a ‘distributed way’, namely, nodes will only need to ‘know’ the weights of their own links and exchange a minimum amount of information with their ‘neighboring’ nodes.

6. Queue length based dynamic strategies. Stability results

Consider the model described in Sections 3 and 4. Without loss of generality—rather with a *gain* of generality—we assume that a transmission attempt (with nonzero probability) is *allowed* on any link i in any slot t , regardless of whether or not there are customers in the queue available for transmission in that slot. (Any transmission attempt interferes with transmissions on the other links in the usual way.) In particular, it is possible to have a ‘successful transmission attempt’ on link i that transmits *no* customer from the corresponding queue, because none were available. We also adopt the convention that the $A_i(t)$ customers arriving at link i in slot t are *not* counted into the queue length $Q_i(t)$ at time t , but *are* immediately available for transmission at time t . (This convention is nonessential, made just to make the expressions ‘cleaner’.) Given our conventions, we obviously have the following recurrence for the queue length:

$$Q_i(t + 1) = [Q_i(t) + A_i(t) - h_i(t)]^+, \quad t = 0, 1, 2, \dots, \tag{6}$$

where $h_i(t) = 1$ if there was a *successful* transmission attempt on link i at time t , and $h_i(t) = 0$ otherwise.

Below we define the two dynamic strategies that we are going to study, which we collectively refer to as the QRA algorithms.

Definition. (*QRA algorithms.*) Each node n maintains *dynamic weights* $w_i(t)$, $t = 0, 1, 2, \dots$, of ‘its’ outgoing links, depending on the corresponding queue lengths. Each node n sets its access probabilities in slot t according to (5). In other words, the set of access probabilities in the network at time t is given by $p(w(t))$. The *QRA-I* algorithm uses weights $w_i(t) = \alpha_i + \gamma_i Q_i(t)^\beta$ with parameters $\alpha_i \geq 0$, $\gamma_i > 0$, and $\beta > 0$. The *QRA-II* algorithm uses weights $w_i(t) = \alpha_i \exp(\gamma_i Q_i(t)^\kappa)$ with parameters $\alpha_i > 0$, $\gamma_i > 0$, and $\kappa \in (0, 1)$.

Given our assumptions on the arrival processes (see Section 4), it is clear that, under both QRA algorithms, the queue length (vector) process $Q(t)$, $t = 0, 1, 2, \dots$, is an irreducible (and aperiodic) Markov chain with countable state space.

Theorem 1. *Suppose that the vector of input rates λ is such that*

$$\lambda < \mu^* \text{ for some } \mu^* \in M. \tag{7}$$

Then,

- (i) *the Markov chain $Q = \{Q(t), t = 0, 1, 2, \dots\}$ is ergodic under the QRA-I policy;*
- (ii) *the Markov chain Q is ergodic under the QRA-II policy when the following additional (exponential moment) assumption on the input flows holds:*

$$E[\exp(a_1 A_i(0))] < \infty \text{ for some } a_1 > 0 \text{ and all } i. \tag{8}$$

Remark 1. If input processes are not i.i.d., condition (8) can be relaxed to (26) (See Subsection 9.1)—the latter condition is the one actually used in the proof.

Remark 2. In essence, the QRA-I algorithm was introduced in [11], in a somewhat different context, where the arrival processes and queue lengths were ‘virtual’, and used to enforce the minimum desired throughputs on the links. However, the queueing stability problem was *not* addressed in [11].

Remark 3. As already discussed in the introduction, the QRA-II algorithm is the ‘random access model’ counterpart of the EXP algorithm [16], [17]. (The latter algorithm applies to a different model.) The queue weights $w_i(t)$ used in the original EXP algorithm have the form

$$w_i(t) = \alpha_i \exp\left(\frac{\gamma_i Q_i(t)}{1 + (\overline{\gamma Q}(t))^{1-\kappa}}\right),$$

where $\overline{\gamma Q}(t) = (1/I) \sum_i \gamma_i Q_i(t)$. The form $w_i(t) = \alpha_i \exp(\gamma_i Q_i(t))^\kappa$ was later used in [9] (again, for a different model). These two forms result in equivalent behavior to FSPs and LFSPs, which are used in our stability proofs, and thus, in principle, either form can be used in QRA-II. The latter form is better suited for a distributed implementation. We emphasize that previous stability analyses in [9], [16], and [17] are for a different model, and do not apply to the model of this paper.

7. Proof of Theorem 1

We prove stability using the fluid limit technique [4], [5], [7], [15], [18]. (For an application of this technique in a discrete-time setting, similar to that of this paper, cf. [2].)

Let $Q^{(r)} = (Q^{(r)}(t), t = 0, 1, 2, \dots)$ denote a queue length process Q with a fixed initial condition such that $\|Q^{(r)}(0)\| = r$, $r > 0$. In the analysis to follow, all variables associated with a process $Q^{(r)}$ will be supplied with the upper index (r) . It will be convenient to extend the definition of the process $Q^{(r)}$ to *continuous time* $t \geq 0$ by adopting the convention that $Q(t) = Q(\lfloor t \rfloor)$.

The following result follows from the state-dependent Lyapunov-type stability criterion for countable Markov chains, first obtained in [13]. (In the specific form (9), below, a Markov chain ergodicity criterion was introduced in [15].)

Proposition 3. *Suppose that there exist constants $\varepsilon > 0$ and $T > 0$ such that, for any sequence of processes $\{Q^{(r)}, r \uparrow \infty\}$, we have*

$$\limsup_{r \rightarrow \infty} \mathbb{E} \left[\frac{1}{r} \|Q^{(r)}(rT)\| \right] \leq 1 - \varepsilon. \tag{9}$$

Then the (original, discrete-time) Markov chain Q is ergodic.

To verify condition (9), the fluid limit technique introduces the sequence of fluid-scaled processes,

$$q^{(r)}(t) := \frac{1}{r} Q^{(r)}(rt), \quad t \geq 0.$$

Note that it suffices to verify (9) under the additional condition that the sequence of rescaled initial states $(1/r)Q^{(r)}(0)$ converges, which is equivalent to

$$q^{(r)}(0) \rightarrow q(0) \quad \text{as } r \rightarrow \infty, \quad \text{where } \|q(0)\| = 1. \tag{10}$$

Then, we establish the following result.

Theorem 2. *There exist constants $\varepsilon > 0$ and $T > 0$ for which the following holds. Consider a fixed sequence of rescaled processes $\{q^{(r)}, r \uparrow \infty\}$, satisfying condition (10). Then, all processes of the sequence can be constructed on a common probability space, such that the following holds with probability 1. Any subsequence of the sequence of realizations of $\{q^{(r)}, r \uparrow \infty\}$ has in turn a further subsequence u.o.c. converging to a trajectory $q = (q(t), t \geq 0)$, which we call an FSP, and which, in particular, satisfies the following properties:*

$$\|q(0)\| = 1, \tag{11}$$

$$\text{function } q(t), t \geq 0, \text{ is Lipschitz continuous,} \tag{12}$$

$$\|q(T)\| \leq 1 - \varepsilon. \tag{13}$$

Theorem 2 will be proved separately for QRA-I (under the assumptions of Theorem 1(i)) and QRA-II (under the assumptions of Theorem 1(ii)) in Sections 8 and 9, respectively. FSPs will be defined differently for QRA-I and QRA-II. We emphasize that not only will the equations describing their dynamics be different, but that their *definitions* are in fact somewhat different as well.

Once Theorem 2 is established (for both QRA-I and QRA-II), this completes the proof of Theorem 1. Indeed, for any fixed $T > 0$, the uniform integrability of the family of random variables $(1/r)\|Q^{(r)}(rT)\|$ (indexed by r , as in Proposition 3) is easily established, using majorization of the queue lengths by the cumulative arrival processes [5], [15]. This fact and Theorem 2 verify the assertion of Proposition 3.

8. Proof of Theorem 2 for QRA-I

8.1. Probability space construction and other preliminaries

Let us denote by

$$F_i^{(r)}(t) := \sum_{s=0}^{t-1} A_i(s), \quad t = 0, 1, 2, \dots, i \in \mathcal{I},$$

the total number of customer arrivals at link i by (and excluding) integer time t . We also denote by

$$\hat{F}_i^{(r)}(t) = \sum_{s=0}^{t-1} h_i(s), \quad t = 0, 1, 2, \dots, i \in \mathcal{I},$$

the total number of successful transmissions on link i by (and excluding) time t .

Without loss of generality, we will assume that the random transmission attempt decision by node n at time t , given its ‘current’ (depending on t) set of access probabilities p_{nm} , $m \in \mathcal{D}_n$, is determined by the random variable $y_n(t)$, uniformly distributed in $[0, 1]$. (Random variables $y_n(t)$ are mutually independent across all n and t .) Namely, node n assumes some fixed ordering of the nodes in \mathcal{D}_n : m_1, \dots, m_ℓ . Then, if $y_n(t) \in \phi_{n,m_1} := [0, p_{n,m_1}]$, node n attempts transmission on link (n, m_1) ; if $y_n(t) \in \phi_{n,m_2} := (p_{n,m_1}, p_{n,m_1} + p_{n,m_2}]$, node n attempts transmission on link (n, m_2) ; and so on. If $y_n(t) \in (p_n, 1]$, node n does not attempt a transmission. We set

$$Y^{(r)}(t, \xi) = \sum_{s=0}^{t-1} I\{y_n(s) \leq \xi_n, n \in \mathcal{N}\}, \quad t = 0, 1, 2, \dots, \xi = (\xi_1, \dots, \xi_N) \in [0, 1]^N, \tag{14}$$

where $I\{\cdot\}$ is an event indicator (not to be confused with I as the number of links).

We will use vector notation, $F^{(r)}(t) = (F_i^{(r)}(t), i \in \mathcal{I})$ and $\hat{F}^{(r)}(t) = (\hat{F}_i^{(r)}(t), i \in \mathcal{I})$. Finally, we extend the time domain of the functions $F^{(r)}(t)$, $\hat{F}^{(r)}(t)$, and $Y^{(r)}(t, \xi)$ to all real $t \geq 0$ by adopting the convention (as we already did for $Q^{(r)}(t)$) that they are constant within each time slot $[\ell, \ell + 1)$ for all integers $\ell \geq 0$.

By our definitions and conventions, $F^{(r)}(0) = 0$, $\hat{F}^{(r)}(0) = 0$, and $Y^{(r)}(0, \xi) = 0$ for every $\xi \in [0, 1]^N$. We also have the following ‘integral form’ of recurrence (6):

$$Q_i^{(r)}(t) = Q_i^{(r)}(0) + F_i^{(r)}(t) - \hat{F}_i^{(r)}(t) - \left[\min_{s \in [0, t]} \{Q_i^{(r)}(0) + F_i^{(r)}(s) - \hat{F}_i^{(r)}(s)\} \right]^-, \quad t \geq 0, i \in \mathcal{I}. \tag{15}$$

Without loss of generality, we can assume that the processes $F^{(r)}(\cdot)$ and $Y^{(r)}(\cdot, \cdot)$, although carrying the index (r) , do not depend on r . (For every r , they are constructed from the same underlying sequences of i.i.d. random variables.) Then, along any fixed sequence of $r \uparrow \infty$, the following functional strong law of large numbers properties hold, with probability 1:

$$\left(\frac{1}{r}\right) F_i^{(r)}(rt) \rightarrow \lambda_i t \quad \text{u.o.c. for all } i \in \mathcal{I}, \tag{16}$$

$$\left(\frac{1}{r}\right) Y^{(r)}(rt, \xi) \rightarrow \left(\prod_{n \in \mathcal{N}} \xi_n\right) t \quad \text{u.o.c.} \tag{17}$$

(In (16) and (17), u.o.c. means that the convergence is uniform on t and (t, ξ) , respectively, within any bounded subset.)

8.2. Definition of FSPs

Consider a sequence of $r \uparrow \infty$. For each r , let $(Q^{(r)}(\cdot), F^{(r)}(\cdot), \hat{F}^{(r)}(\cdot), Y^{(r)}(\cdot, \cdot))$ be a realization (that is, a fixed sample path) of the corresponding random process, with some fixed initial condition $Q^{(r)}(0)$, $\|Q^{(r)}(0)\| = r$. The entire realization is uniquely determined

by $Q^{(r)}(0)$, $F^{(r)}(\cdot)$, and $Y^{(r)}(\cdot, \cdot)$. Assume that this sequence of realizations satisfies conditions (16) and (17).

Consider the following rescaled trajectory for each r :

$$(q^{(r)} = (q^{(r)}(t), t \geq 0), f^{(r)} = (f^{(r)}(t), t \geq 0), \hat{f}^{(r)} = (\hat{f}^{(r)}(t), t \geq 0)),$$

where $f^{(r)}(t) = (1/r)F^{(r)}(rt)$, $\hat{f}^{(r)}(t) = (1/r)\hat{F}^{(r)}(rt)$, and (recall $q^{(r)}(t) = (1/r)Q^{(r)}(rt)$).

A triple of vector functions $(q = (q(t), t \geq 0), f = (f(t), t \geq 0), \hat{f} = (\hat{f}(t), t \geq 0))$ is called an FSP if the u.o.c. convergence,

$$(q^{(r)}, f^{(r)}, \hat{f}^{(r)}) \rightarrow (q, f, \hat{f}), \tag{18}$$

holds for at least one sequence $(q^{(r)}, f^{(r)}, \hat{f}^{(r)})$ of scaled trajectories (with $r \uparrow \infty$), such that (16) and (17) hold.

We note that the definition of the FSPs given here, as well as their dynamic properties derived in Subsection 8.3, below, do *not* require condition (7).

8.3. Basic dynamics of FSPs

Lemma 1. Any FSP satisfies the following conditions: (11),

$$\text{functions } q(\cdot), f(\cdot), \text{ and } \hat{f}(\cdot), \text{ are Lipschitz continuous with } f_i(t) = \lambda_i t, \tag{19}$$

$$q'_i(t) = [\lambda_i - v_i(t)]^+_i \text{ for all } i \in \mathcal{I} \text{ and for almost all } t \geq 0, \tag{20}$$

$$v(t) \in \arg \max_{x \in M} \{\gamma \times q(t)^\beta\} \log x \tag{21}$$

Proof. Consider a fixed FSP, and any fixed sequence of rescaled paths $(q^{(r)}, f^{(r)}, \hat{f}^{(r)})$ which ‘defines’ this FSP (via convergence (18) and the other conditions). Properties (11) and (19) are obvious, given the FSP construction. (The property that $f_i(t) = \lambda_i t$ follows from (16).) Therefore, almost all points $t \geq 0$ are such that proper derivatives of component functions of the FSP exist—such time points will be called *regular*. It will suffice to show (20) and (21) for a given regular point $t > 0$. Indeed, switching to rescaled paths in (15) and taking the limit on r , we obtain

$$q_i(t) = q_i(0) + f_i(t) - \hat{f}_i(t) - \left[\min_{s \in [0, t]} \{q_i(0) + f_i(s) - \hat{f}_i(s)\} \right]^-, \quad t \geq 0, i \in \mathcal{I}. \tag{22}$$

Set $v_i(t) = \hat{f}'_i(t)$, and recall that $f'_i(t) = \lambda_i$. Given (22), we see that the equation in (20) holds trivially if $q_i(t) > 0$. If $q_i(t) = 0$, we must have $q'_i(t) = 0$ (by regularity), and then $\lambda_i - v_i(t) \leq 0$, because otherwise (22) would imply that the right derivative $(d^+/dt)q_i(t)$ exists and is equal to $\lambda_i - v_i(t) > 0$. This proves (20).

Let us prove (21). Now, consider the FSP, and the corresponding rescaled trajectories in a small interval $[t, t + \varepsilon]$. By continuity of $q(\cdot)$ and by (18), all the values of $q^{(r)}(s) \equiv Q^{(r)}(rs)$, $s \in [t, t + \varepsilon]$, for all sufficiently large r are close to $q(t)$ as long as $\varepsilon > 0$ is small. For each r , consider the corresponding unscaled trajectory of $Q^{(r)}(\cdot)$ in the (corresponding) time interval $[rt, rt + r\varepsilon]$. It will be convenient to assume that the trajectory with index r uses rescaled dynamic weights $w_i^{(r)}(s) = w_i(s)/r = \alpha_i/r + \gamma_i(Q_i^{(r)}(s))^\beta/r$. (This does not change anything, by Proposition 2(a).) We see that, in $[rt, rt + r\varepsilon]$, $w_i^{(r)}(s)$ is close to $\gamma_i q_i(t)^\beta$. Then, it is easy to see from Proposition 2(c) that, for any link $i = (n, m)$ with $q_i(t) > 0$, the corresponding interval $\phi_{n,m}$ (see the construction governing transmission attempts) is close to

the analogous interval corresponding to the access probabilities $p(\gamma \times q(t)^\beta)$. (Again, this is for the unscaled trajectory in the interval $[rt, rt + r\varepsilon]$ with small ε and for large values of r .) Given (14), condition (17), and again Proposition 2(c), this implies that the time average rate of successful attempts in the interval $[rt, rt + r\varepsilon]$ is close to $\mu_i(p(\gamma \times q(t)^\beta))$. This shows (we omit ε - δ formalities) that, for all links i with $q_i(t) > 0$, $v_i(t) = \mu_i(p(\gamma \times q(t)^\beta))$, which, by Proposition 2(b), implies (21).

8.4. Stability of FSPs. Conclusion of the proof of Theorem 2

All the statements of Theorem 2, except (13), easily follow from the fact that (16) and (17) hold with probability 1, the FSP construction, and the FSP property, (19). It remains to show that FSPs satisfy (13), which is done in the following theorem.

Theorem 3. *Consider the function*

$$\Psi(y) = \frac{1}{\beta + 1} \sum_i \frac{\gamma_i}{\mu_i^*} y_i^{\beta+1}, \quad y \in \mathbb{R}_+^I.$$

(It depends on μ^* as a parameter.) Suppose that λ satisfies (7), and consider the family of Lipschitz continuous trajectories $(q(t), t \geq 0)$ satisfying (20) and (21). Then, for any $\varepsilon_1 > 0$, there exists $\varepsilon_2 > 0$ such that

$$\Psi(q(t)) \geq \varepsilon_1 \text{ implies that } \frac{d}{dt} \Psi(q(t)) \leq -\varepsilon_2.$$

As a corollary, there exist $\varepsilon > 0$ and $T > 0$ such that (13) holds uniformly for all FSPs.

Remark. Lipschitz continuous trajectories satisfying (20) and (21) also arise in a completely different setting, namely, in the ‘session level’ stability analysis of communication networks with ‘concave-utility-based’ allocation of service rates to different sessions; see [3], [8], and [22]. (We want to emphasize that the *derivation* of properties (20) and (21) in our setting is completely different.) In all the previous work cited above the region M (which has a different meaning from ours) is convex. In our model region M is not necessarily convex, and in fact nonconvex in most cases. However, for the purposes of establishing trajectory stability in Theorem 3, convexity of M is *not* important. Since this last point may not be immediately clear from the previous work (some of which does use convexity of M , even though it does not have to), and for completeness, we present the proof of Theorem 3.

Proof of Theorem 3. Let $\log \mu^* = u^*$ and $\log v(t) = u(t)$. Then, we have

$$\begin{aligned} \frac{d}{dt} \Psi(q(t)) &= \sum_i \frac{\gamma_i}{\mu_i^*} q_i(t)^\beta q_i'(t) \\ &= \sum_i \frac{\gamma_i}{\mu_i^*} q_i(t)^\beta (\lambda_i - \exp(u_i(t))) \\ &\leq \sum_i \frac{\gamma_i}{\mu_i^*} q_i(t)^\beta (\lambda_i - \exp(u_i^*) - \exp(u_i^*)(u_i(t) - u_i^*)) \end{aligned} \tag{23}$$

$$= \sum_i \gamma_i q_i(t)^\beta \left(\frac{\lambda_i}{\mu_i^*} - 1 \right) - \left(\sum_i \gamma_i q_i(t)^\beta u_i(t) - \sum_i \gamma_i q_i(t)^\beta u_i^* \right) \tag{24}$$

$$\leq \sum_i \gamma_i q_i(t)^\beta \left(\frac{\lambda_i}{\mu_i^*} - 1 \right). \tag{25}$$

Inequality (23) uses convexity of the exponent function, and (24) is less than or equal to (25) because $\sum_i \gamma_i q_i(t)^\beta u_i(t) - \sum_i \gamma_i q_i(t)^\beta u_i^* \geq 0$, by condition (21). Expression (25) is negative and bounded away from 0 as long as $\Psi(q(t))$ is positive and bounded away from 0.

9. Proof of Theorem 2 for QRA-II

9.1. Preliminaries

Recall that we are now in the conditions of Theorem 1(ii), and, therefore, we are under the additional assumption (8), which implies a large deviations (Cramer’s) bound for the input processes. For any $i \in \mathcal{I}$ and any $\nu > 0$, there exists a constant $a = a(\nu) > 0$ such that, for all sufficiently large n , uniformly on $k \geq 1$,

$$P \left\{ \left| \frac{1}{n} \sum_{t=k}^{k+n-1} A_i(t) - \lambda_i \right| \geq \nu \right\} < e^{-an}. \tag{26}$$

Let arbitrary $\nu > 0$ and $L > 0$ be fixed. Let us also pick any $\zeta > 0$ such that $\zeta < \eta \equiv 1 - \kappa$. For each n , let us cover the interval $[0, rL]$ with $P_L^r := \lfloor rL/r^\zeta \rfloor + 1$ equal nonoverlapping r^ζ -long intervals $[(j - 1)r^\zeta, jr^\zeta]$, $1 \leq j \leq P_L^r$. For each $i \in \mathcal{I}$ and each $\xi \in [0, 1]$, define the number of arrivals of flow i in the time interval $[(j - 1)r^\zeta, jr^\zeta]$ as

$$F_{i,j}^{(r)} := F_i^{(r)}(jr^\zeta) - F_i^{(r)}((j - 1)r^\zeta)$$

and define

$$Y_{i,j,\xi}^{(r)} := Y_i^{(r)}(jr^\zeta, \xi) - Y_i^{(r)}((j - 1)r^\zeta, \xi).$$

Let

$$E_i^r(L, \nu) = \bigcup_{1 \leq j \leq P_L^r} \left(\left| \frac{F_{i,j}^{(r)}}{r^\zeta} - \lambda_i \right| > \nu \right),$$

$$G_i^r(L, \nu, \xi) = \bigcup_{1 \leq j \leq P_L^r} \left(\left| \frac{Y_{i,j,\xi}^{(r)}}{r^\zeta} - \xi \right| > \nu \right).$$

Lemma 2. *The following properties hold (with \mathcal{Q}_+ being the set of strictly positive rational numbers):*

$$P \left\{ \bigcup_{\nu, L \in \mathcal{Q}_+} \bigcap_{k=1}^\infty \bigcup_{r=k}^\infty E_i^r(L, \nu) \right\} = 0 \quad \text{for all } i \in \mathcal{I}, \tag{27}$$

$$P \left\{ \bigcup_{\nu, L \in \mathcal{Q}_+} \bigcup_{\xi \in [0, 1], \xi \in \mathcal{Q}_+} \bigcap_{k=1}^\infty \bigcup_{r=k}^\infty G_i^r(L, \nu, \xi) \right\} = 0 \quad \text{for all } i \in \mathcal{I}. \tag{28}$$

Equivalently, with probability 1, for any rational numbers $L > 0$, $\nu > 0$, and $\xi \in [0, 1]$, there exists finite k such that, for all $r > k$,

$$\max_{i \in \mathcal{I}, 1 \leq j \leq P_L^r} \left| \frac{E_i^r(L, \nu)}{r^\zeta} - \lambda_i \right| \leq \nu, \tag{29}$$

$$\max_{i \in \mathcal{I}, 1 \leq j \leq P_L^r} \left| \frac{G_i^r(L, \nu, \xi)}{r^\zeta} - \xi \right| \leq \nu. \tag{30}$$

Proof. The proof uses Cramer’s bound (26) and Borel–Cantelli’s lemma. (See [16, Lemma 4.3] for the proof of (27); (28) is proved analogously.)

9.2. FSP definition and stability. Proof of Theorem 2

Consider the system under the QRA-II algorithm. For this algorithm, we define FSPs in exactly the same way as for QRA-I (see Subsection 8.2), except that we require, in addition to (or, rather, instead of) (16) and (17), a defining sequence $(q^{(r)}, f^{(r)}, \hat{f}^{(r)})$ of scaled sample paths that satisfies the stronger conditions given in (29) and (30). (Such FSPs clearly satisfy the initial condition, (11), and the Lipschitz condition, (19), but certainly *not* conditions (20) and (21)—the FSP dynamics under QRA-II is very different.) With this FSP definition, the proof of all the statements of Theorem 2, except (13), is almost automatic. To establish (13), we prove the following key property of the FSPs under the QRA-II algorithm.

Theorem 4. *Suppose that a point $\mu^* \in M \cap \mathbb{R}_{++}^I$ is such that, for some real c , $\gamma_i(\mu_i^* - \lambda) = c$ for all i . (We do not necessarily assume (7) and, thus, a positive c may or may not exist.) Then, in addition to (11) and (12), every FSP satisfies the following condition at every regular point $t \geq 0$, where $\max_i \gamma_i q_i(t) > 0$:*

$$\frac{d}{dt} \max_i \gamma_i q_i(t) \leq -c. \tag{31}$$

The above definition of the FSPs for QRA-II, as well as that for QRA-I, does *not* require condition (7). And Theorem 4 holds regardless of (7). If condition (7) does hold then μ^* can be chosen to be the point where the ray in the direction $(1/\gamma_1, \dots, 1/\gamma_I)$, starting at λ , hits the boundary of M , in which case $c > 0$ and

$$\frac{d}{dt} \max_i \gamma_i q_i(t) \leq -c < 0,$$

which of course proves (13).

Proof of Theorem 4. Properties (11) and (12) are proved in the same way as for the QRA-I algorithm (and for fluid limits in general). To prove (31), consider a fixed FSP and a sequence of rescaled prelimit trajectories defining it. We will use the notation $z(t) := \max_i \gamma_i q_i(t)$, $z^{(r)}(t) := \max_i \gamma_i q_i^{(r)}(t)$, and $Z^{(r)}(rt) := \max_i \gamma_i Q_i^{(r)}(rt)$. Suppose that (31) does not hold. Then, there exists a regular point t such that $z(t) > 0$ and $z'(t) > -c_1 > -c$. We will show that this leads to a contradiction. We can choose constants $\delta > 0$, $\delta_1 > 0$, and $c_2 \in (c_1, c)$ such that

$$z(s) > \delta_1 \quad \text{for all } s \in [t, t + \delta]$$

and

$$\frac{z(t + \delta) - z(t)}{\delta} > -c_2.$$

For each r , let us now divide the interval $[t, t + \delta]$ into $r^\kappa \delta / \ell$ intervals, each of length $\ell r^\eta / r$, where $\eta = 1 - \kappa$ and $\ell > 0$ is an arbitrary fixed constant. (Since $r^\kappa \delta / \ell$ may not be an integer, we should divide the interval into, say, $\lceil r^\kappa \delta / \ell \rceil$ intervals. To avoid trivial complications and heavy notation, we assume that $r^\kappa \delta / \ell$ is an integer. It will be clear that we do not lose the correctness of the argument.) Note that in the ‘unscaled time’ each subinterval is of length ℓr^η .

From the Dirichlet principle, for all sufficiently large r , in at least one of the subintervals (of length $\ell r^\eta / r$), the average rate of change of $z^{(r)}(\cdot)$ is greater than or equal to $-c_2$. We pick such a subinterval $[s^{(r)}, s^{(r)} + \ell r^\eta / r]$ for each r . Let us choose a further subsequence of the sequence of indices $\{r\}$ (which we will still denote by $\{r\}$) such that $s^{(r)} \rightarrow s$ for some fixed $s \in [t, t + \delta]$. Obviously, the right endpoint $s^{(r)} + \ell r^\eta / r$ of the subinterval also converges to s .

From the subsequence $\{r\}$, we choose a further subsequence such that the order of values of $\gamma_i q_i^{(r)}(s^{(r)})$, $i \in \mathcal{I}$, remains the same. For example, without loss of generality, we can assume that

$$z^{(r)}(s^{(r)}) = \gamma_1 q_1^{(r)}(s^{(r)}) \geq \dots \geq \gamma_I q_I^{(r)}(s^{(r)}).$$

Finally, for each $i \in \mathcal{I}$, consider the following trajectories:

$$y_i^{(r)}(\tau) := r^\kappa \left(\gamma_i q_i^{(r)} \left(s^{(r)} + \frac{\tau}{r^\kappa} \right) - z^{(r)}(s^{(r)}) \right), \quad \tau \in [0, \ell],$$

and choose a subsequence such that, for each i ,

$$y_i^{(r)}(0) \rightarrow y_i(0),$$

where $\max_i y_i(0) = y_1(0) = 0$ (by our construction) and each other $y_i(0)$ is either finite nonpositive or $-\infty$. Let us consider only the case when all the $y_i(0)$ are finite. (If not, it is easy to observe that, in the (unscaled) time interval $[rs^{(r)}, rs^{(r)} + r^\eta \ell]$, the queues with $y_i(0) = -\infty$ have asymptotically vanishing impact on the service of queues with $y_i(0) > -\infty$. So, essentially the same argument, restricted to the subset of queues with finite $y_i(0)$, applies.)

We note that the trajectory $y_i^{(r)}(\cdot)$ is obtained from the trajectory $Q_i^{(r)}(\cdot)$ by the time ‘speedup’ of r^η and the ‘space’ scaling by the factor $1/r^\eta$ (in addition to the time shift).

In the next step we can and do choose a further subsequence of $\{r\}$ such that the sequence $\{y_i^{(r)}(\cdot), i \in \mathcal{I}\}$, converges u.o.c. in $[0, \ell]$ to a Lipschitz continuous trajectory $\{y_i(\cdot), i \in \mathcal{I}\}$, which we call an LFSP. In this step we use properties (29) and (30) which guarantee that, roughly speaking, the functional law of large numbers holds not only over any (unscaled) interval of length δr , but also uniformly over the set of ℓr^η -long subintervals of it.

It is also not hard to see that the LFSP trajectory satisfies conditions (32) and (33), below, at every regular point $\tau \in [0, \ell]$.

$$y_i'(\tau) = \gamma_i(\lambda_i - \mu_i(\tau)) \quad \text{for all } i \in \mathcal{I}, \tag{32}$$

$$\mu_i(\tau) \in \arg \max_{x \in M} \{\alpha \times \exp(a_2 y(\tau))\} \log x, \tag{33}$$

where $a_2 = \kappa/z(s)^\eta$. Indeed, note that $z^{(r)}(s^{(r)})/z(s) \rightarrow 1$ and $\gamma_i q_i^{(r)}(s^{(r)} + \tau/r^\kappa)/z(s) \rightarrow 1$ as $r \rightarrow \infty$, uniformly in i and in $\tau \in [0, \ell]$; recall that $\gamma_i q_i^{(r)}(s^{(r)} + \tau/r^\kappa) - z^{(r)}(s^{(r)}) = r^{-\kappa} y_i^{(r)}(\tau)$. We have

$$\begin{aligned} (\gamma_i Q_i^{(r)}(s^{(r)}r + r^\eta \tau))^\kappa - Z^{(r)}(s^{(r)}r)^\kappa &= r^\kappa \left(\left(\gamma_i q_i^{(r)} \left(s^{(r)} + \frac{\tau}{r^\kappa} \right) \right)^\kappa - z^{(r)}(s^{(r)})^\kappa \right) \\ &= r^\kappa (\kappa \hat{z}^{\kappa-1}) (r^{-\kappa} y_i^{(r)}(\tau)), \end{aligned}$$

where \hat{z} is a number ‘between’ $z^{(r)}(s^{(r)})$ and $\gamma_i q_i^{(r)}(s^{(r)} + \tau/r^\kappa)$. (Here we simply applied the mean value theorem for an increment of function x^κ .) But, $\hat{z} \rightarrow z(s)$ as $r \rightarrow \infty$, uniformly in i and τ , and, therefore, we have the following uniform convergence:

$$(\gamma_i Q_i^{(r)}(s^{(r)}r + r^\eta \tau))^\kappa - Z^{(r)}(s^{(r)}r)^\kappa \rightarrow \frac{\kappa}{z(s)^\eta} y_i(\tau).$$

It remains to consider the behavior of the sample path in the unscaled time interval $[s^{(r)}r + r^\eta \tau, s^{(r)}r + r^\eta \tau + r^\eta \Delta \tau]$ with small fixed $\Delta \tau$, as r becomes large, and use the form of the

QRA-II scheduling rule. (The argument here is analogous to the one used to prove (21) for the QRA-I rule. We omit details.)

By definition of an LFSP, $\max_i y_i(0) = 0$, and, by our construction,

$$\max_i y_i(\ell) \geq -c_2\ell > -c\ell. \tag{34}$$

We will use the following Lyapunov function:

$$\Psi(\tau) = \sum_i \frac{1}{\mu_i^* \gamma_i} \alpha_i \exp(a_2(y_i(\tau) + c\tau)),$$

where, recall, $\mu^* \in M$ is such that $\gamma_i(\mu_i^* - \lambda_i) = c$ for all i . We have

$$\Psi(0) \leq \sum_i \frac{\alpha_i}{\mu_i^* \gamma_i}.$$

We let $u_i(\tau) = \log \mu_i(\tau)$ and $u_i^* = \log \mu_i^*$. Then,

$$\begin{aligned} \frac{d}{dt} \Psi(y(\tau)) &= a_2 \sum_i \frac{\alpha_i}{\mu_i^* \gamma_i} \exp(a_2(y_i(\tau) + c\tau))(y_i'(\tau) + c) \\ &= a_2 \sum_i \frac{\alpha_i}{\mu_i^*} \exp(a_2(y_i(\tau) + c\tau))(\mu_i^* - \mu_i(\tau)) \\ &\leq a_2 \sum_i \frac{\alpha_i}{\mu_i^*} \exp(a_2(y_i(\tau) + c\tau))(\mu_i^* - (\mu_i^* + \mu_i^*(u_i(\tau) - u_i^*))) \\ &= a_2 \exp(a_2c\tau) \left(\sum_i \alpha_i \exp(a_2y_i(\tau)) u_i^* - \sum_i \alpha_i \exp(a_2y_i(\tau)) u_i(\tau) \right) \\ &\leq 0. \end{aligned}$$

Since $\Psi(\tau)$ is nonincreasing and $\Psi(0)$ is bounded, we see that, for some constant K , depending only on the system parameters and on c ,

$$\max_i y_i(\tau) \leq Kz(s)^\eta - c\tau, \quad \tau \in [0, \ell]. \tag{35}$$

Recall that the function $z(\cdot)$ is defined on the time scale of the FSP (not the LFSP), and it is Lipschitz and $z(0) \leq \max \gamma_i$. Then $Kz(s)^\eta$ is uniformly bounded within any finite interval (of the FSP time scale) and, in particular, for $s \in [t, t + \delta]$, with t and δ chosen at the beginning of this proof. Therefore, since ℓ could be chosen arbitrarily large, (35) contradicts (34).

References

- [1] ANANTHARAM, V. (1991). The stability region of the finite-user slotted Aloha protocol. *IEEE Trans. Inf. Theory* **37**, 535–540.
- [2] ANDREWS, M. *et al.* (2004). Scheduling in a queueing system with asynchronously varying service rates. *Prob. Eng. Inf. Sci.* **18**, 191–217.
- [3] BONALD, T. AND MASSOULIE, L. (2001). Impact of fairness on Internet performance. In *Proc. ACM SIGMETRICS*, ACM, New York, pp. 82–91.
- [4] CHEN, H. (1995). Fluid approximations and stability of multiclass queueing networks: work-conserving disciplines. *Ann. Appl. Prob.* **5**, 637–665.

- [5] DAI, J. G. (1995). On the positive Harris recurrence for open multiclass queueing networks: a unified approach via fluid limit models. *Ann. Appl. Prob.* **5**, 49–77.
- [6] DAI, J. G. AND LIN, W. (2005). Maximum pressure policies in stochastic processing networks. *Operat. Res.* **53**, 197–218.
- [7] DAI, J. G. AND MEYN, S. P. (1995). Stability and convergence of moments for open multiclass queueing networks via fluid limit models. *IEEE Trans. Automatic Control* **40**, 1889–1904.
- [8] DE VECIANA, G., LEE, T. J., AND KONSTANTOPOULOS, T. (2001). Stability and performance analysis of networks supporting elastic services. *IEEE/ACM Trans. Networking* **9**, 2–14.
- [9] ERYILMAZ, A., SRIKANT, R. AND PERKINS, J. (2005). Stable scheduling policies for fading wireless channels. *IEEE/ACM Trans. Networking* **13**, 411–424.
- [10] GUPTA, P. AND STOLYAR, A. L. (2005). Throughput region of random access networks of general topology. Submitted.
- [11] GUPTA, P. AND STOLYAR, A. L. (2006). Optimal throughput allocation in general random-access networks. In *Proc. 40th Annual Conf. Inf. Sci. Systems* (Princeton, NJ, March 2006), IEEE, pp. 1254–1259.
- [12] KAR, K., SARKAR, S. AND TASSIULAS, L. (2004). Achieving proportionally fair rates using local information in Aloha networks. *IEEE Trans. Automatic Control* **49**, 1858–1862.
- [13] MALYSHEV, V. A. AND MENSHIKOV, M. V. (1979). Ergodicity, continuity, and analyticity of countable Markov chains. *Trans. Moscow Math. Soc.* **39**, 3–48.
- [14] MASSEY, J. AND MATHYS, P. (1985). The collision channel without feedback. *IEEE Trans. Inf. Theory* **31**, 192–204.
- [15] RYBKO, A. N. AND STOLYAR, A. L. (1992). Ergodicity of stochastic processes describing the operation of open queueing networks. *Problems Inf. Transmission* **28**, 199–220.
- [16] SHAKKOTTAI, S. AND STOLYAR, A. L. (2002). Scheduling for multiple flows sharing a time-varying channel: the exponential rule. In *Analytic Methods in Applied Probability* (Amer. Math. Soc. Transl. Ser. 2 **207**), American Mathematical Society, Providence, RI, pp. 185–201.
- [17] SHAKKOTTAI, S., SRIKANT, R. AND STOLYAR, A. L. (2004). Pathwise optimality of the exponential scheduling rule for wireless channels. *Adv. Appl. Prob.* **36**, 1021–1045.
- [18] STOLYAR, A. L. (1995). On the stability of multiclass queueing networks: a relaxed sufficient condition via limiting fluid processes. *Markov Process. Relat. Fields* **1**, 491–512.
- [19] STOLYAR, A. L. (2004). Max-weight scheduling in a generalized switch: state space collapse and workload minimization in heavy traffic. *Ann. Appl. Prob.* **14**, 1–53.
- [20] TASSIULAS, L. AND EPHREMIDES, A. (1992). Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Trans. Automatic Control* **37**, 1936–1948.
- [21] WANG, X. AND KAR, K. (2004). Distributed algorithms for max-min fair rate allocation in Aloha networks. In *Proc. 42nd Annual Allerton Conf.* (Urbana-Champaign, October 2004).
- [22] YE, H. Q. (2003). Stability of data networks under an optimization-based bandwidth allocation. *IEEE Trans. Automatic Control* **48**, 1238–1242.