Queue Back-Pressure Random Access in Multi-Hop Wireless Networks: Optimality and Stability

Jiaping Liu†, Alexander L. Stolyar‡, Mung Chiang† and H. Vincent Poor †

Abstract—**A model for wireless networks with slotted-Aloha-type random access and with multi-hop flow routes is considered. The goal is to devise distributed algorithms for utility-optimal end-to-end throughput allocation and queueing stability. A class of *queue back-pressure random access algorithms (QBRAs)*, in which actual queue lengths of the flows in each node's close neighborhood are used to determine the nodes' channel access probabilities, is studied. This is in contrast to some previously proposed algorithms, which are based on deterministic optimization formulations and are oblivious to actual queues. QBRA is also substantially different from the well studied "MaxWeight" type scheduling algorithms, even though both use the concept of back-pressure.**

For the model with infinite backlog at each flow source, it is shown that QBRA, combined with simple congestion control local to each source, leads to optimal end-to-end throughput allocation within the network *saturation throughput region* achievable by random access, without end-to-end message passing. This scheme is generalized to the case with minimum flow rate constraints. For the model with stochastic exogenous arrivals, it is shown that QBRA ensures stability of the queues as long as nominal loads of the nodes are within the saturation throughput region. Simulation comparison of QBRA and the queue oblivious random access algorithms, shows that QBRA reduces end-to-end delays.

Index Terms

Aloha, Random Access, Distributed Algorithm, Queue Back-Pressure, Stability, Throughput Region.

I. Introduction

In wireless ad hoc networks, contention resolution and interference management among links are among the most important issues, which motivates the extensive study of wireless medium access control (MAC) protocols. The standard MAC protocol currently used in IEEE 802.11 [3] is the Distributed Coordination Function (DCF) with Binary Exponential Backoff (BEB) mechanism. However, it has been concluded by many researchers that DCF with BEB mechanism for contention control can be inefficient and unfair, eg. [11]. Thus, there are significant challenges in designing MAC protocols that are both efficient in terms of throughput, latency, energy consumption, etc., and allow distributed implementation minimizing signalling or message passing overhead.

It has been shown that the maximum throughput region can be achieved by much studied "MaxWeight"-type scheduling algorithms as originally proposed in [14]. However, in the context of wireless networks, MaxWeight algorithms typically need to be centralized for implementation. Some recent works (e.g. see [8], [12] and the references therein) propose distributed solutions of the MaxWeight algorithm, but such implementations also require heavy signaling procedures whose complexity relies on the size of the network. The impact of central coordination, and excessive signalling overhead on the overall performance has not been quantified, or extensively studied, and yet remains unclear. Another class of wireless scheduling schemes, known as random access ("slotted-Aloha-type") algorithms, typically provide smaller throughput regions, but are simpler and more amenable to distributed implementations. In this paper we consider a model of random access for multi-hop transmissions.

Random access models have been widely adopted in contemporary works, such as [1], [2], [4]–[7], [9], [10], [13], [15] and [16]. Informally, we can classify them into two categories: "pure optimization-based" algorithms (e.g. [5], [6], [10], [15] and [9]) and dynamic, queue-length based strategies (e.g. [2], [7], [13] and [4]). Algorithms of the former type solve an optimization problem that allocates network resources (e.g. effective link throughputs) to satisfy and/or optimize traffic demands of different flows; they require optimization parameters to be specified a priori and are typically oblivious to the dynamics of actual queues in the network. Moreover, the iterative algorithms in [5], [15] and [6] involve end-to-end message passing within the network; the revised algorithms proposed in [10] and [9] reduce the signaling to a cluster of interfering nodes but the convergence and optimality have been shown only in a single-hop transmission model. The latter type algorithms, including the Queue-length based Random Access (QRA) algorithm in [2] and [13], and the constant-time distributed scheduling policy which coincides with a certain type of QRA in particular systems in [7] and [4], are operated by adaptively responding to actual queueing dynamics and thus guarantee queueing stability of the system. In particular, for the class of QRA algorithms, even though in some cases they appear to have the same optimization objective as the former optimization-based algorithms, they do not need a priori knowledge of traffic flow input rates to achieve queueing stability if such is feasible. It worth noting that the previous studies of dynamic random access in the latter type ([2], [7], [13] and [4]) all assume a traffic flow model with single-hop transmission without multi-hop routes.

In this paper we propose and study a class of *queue back-pressure random access* (QBRA) algorithms for a multi-hop network, and generally with multi-hop end-to-end flows. The algorithms use flow queue differentials on the links to determine link access probabilities, while the MaxWeight algorithms use queue differentials as well, but in a completely

This work was supported in part by the DARPA CBMANET project, and in part by the National Science Foundation under Grant CNS-06-25637. It was presented in part at the 45th Annual Allerton Conference on Communication, Control, and Computing, Monticello, IL, Sep 2007.

†Department of Electrical Engineering, Princeton University, Princeton, NJ, US. {jiapingl, chiangm, poor}@princeton.edu

‡Bell Labs, Alcatel-Lucent, Murray Hill, NJ 07974, US. stolyar@research.bell-labs.com

different way. Our main contributions are as follows.

(i) For the problem of utility-optimal end-to-end throughput allocation in the model of saturated sources, as considered in Section V, we prove that QBRA combined with extremely simple congestion control at each flow source, solves the problem of *weighted proportional fair* (sum-log utility) end-to-end throughput allocation among the flows. We also prove an extension of this result for the case of additional minimum flow-rate constraints. This generalizes and considerably strengthens the corresponding result in [2]. The result in [2] applies to single-hop flows and proves optimality of equilibrium but not convergence towards equilibrium: it does not state the convergence to an optimal point – only the fact that *if* convergence takes place, then optimality holds. A further generalization - to more general utility functions - is also possible, and will be considered as future work.

(ii) For the problem of queueing stability in the model of exogenous arrivals as considered in Section VI, we prove that QBRA "automatically" ensures stability without knowing input rates, as long as nominal link loads are within the network *saturation throughput region*. This generalizes some of the stability results in [13], which apply to single-hop flows. The stability proof in this paper is conducted with fluid limit techniques but using a novel Lyapunov function that is substantially different than that in [13] - the proof in [13] does not generalize to the multi-hop case. We will elaborate on this in Section V-D.

(iii) Finally, we present simulation results as considered in Section VII, with a variety of parameter setting, showing good performance of QBRA, in particular in terms of end-to-end delays.

II. Basic Notation and Definitions

Typically, we use bold letters $\boldsymbol{x}, \boldsymbol{y}, \ldots$ to denote vectors, as opposed to scalars $x, y, \ldots$. We use the notations $\mathbb{R}$, $\mathbb{R}_+$ and $\mathbb{R}_{++}$ for the set of real, real non-negative and real positive numbers, respectively. Correspondingly, d-times product spaces are denoted as $\mathbb{R}^d$, $\mathbb{R}^d_+$ and $\mathbb{R}^d_{++}$. We write $\boldsymbol{x} \cdot \boldsymbol{y}$ to denote scalar product, and $\|\boldsymbol{x}\| = \sqrt{\boldsymbol{x} \cdot \boldsymbol{x}}$ for the Euclidean norm, inducing the standard metric. Cardinality (i.e. the number of elements) of a finite set $\mathcal{A}$ is denoted by $|\mathcal{A}|$. We denote $[z]^+ = \max\{z, 0\}$.

We use $\prec, \preceq, \succ, \succeq$ for componentwise vector inequalities, e.g. $\boldsymbol{x} \succ \boldsymbol{y}$ means $x_i > y_i, \forall i$. For any scalar function $T : \mathbb{R} \to \mathbb{R}$, $T(\boldsymbol{x}) = (T(x_1), \cdots, T(x_d))$ and for any subset $\mathcal{C} \in \mathbb{R}^d$, $T(\mathcal{C}) = \{T(\boldsymbol{v}) : \boldsymbol{v} \in \mathcal{C}\}$.

III. System Model

A. Wireless Network Model

We consider a wireless multi-hop network described as a directed graph $G = (\mathcal{N}, \mathcal{L})$, where $\mathcal{N}$ is the set of nodes and $\mathcal{L}$ is the set of the logical (directed) communication links between pairs of nodes; t_l and r_l are the transmitter and receiver nodes respectively of link l. There is a finite number of traffic flows, indexed by $r \in \mathcal{R}$; each flow has fixed source and destination nodes, and a fixed route. Here throughout the paper, we will use the terms flow and route interchangeably, and use index r for either one. Let $\mathcal{L}_r \subseteq \mathcal{L}$ denote the set of links on route r, and index links $l \in \mathcal{L}_r$ from source to destination in an ascending order as $l(r, j),\ j = 1, 2, 3, \cdots$. We also assume each node maintains separate queues of data packets of different flows. Let $Q_l^{(r)}$ denote the queue length of flow r packets located in the transmitter node t_l of link $l \in \mathcal{L}_r$. To simplify notation, we often write $Q_j^{(r)}$ to mean $Q_{l(r,j)}^{(r)}$, i.e. for the queue length of flow r at the j-th node in its route.

The system operates in discrete (or, slotted) time $t = 0, 1, 2, \cdots$. In any time slot, each node may attempt to transmit one packet at most on one of its outgoing links. A packet transmission attempt on a link is successful if it is not "interfered with" by another simultaneous transmission during the same time slot; otherwise the transmission fails. The interference model is the same as in [13] and [2], and is somewhat more general than in [15]. First, any transmission attempt to a node will fail if this node is transmitting. Second, if there are two or more simultaneous transmissions to the same node, they all fail. Third, for each node n there is the set of nodes $\mathcal{N}_n \subseteq \mathcal{N}$ with which it interferes, namely, a transmission to any node in $\mathcal{N}_n$ will fail if node n transmits. Note that according to our interference model, $n \in \mathcal{N}_n$ and $\mathcal{D}_n \subseteq \mathcal{N}_n$, where $\mathcal{D}_n \subseteq \mathcal{N} \backslash n$ is the set of nodes m to which node n has data to send. In summary, a transmission attempt on link $l \in \mathcal{L}$ is successful if and only if no node in the set $\{n : n \neq t_l, r_l \in \mathcal{N}_n\}$ transmits.

For each n let us define $\mathcal{S}_n = \{l \in \mathcal{L} : r_l \in \mathcal{N}_n\}$. This set includes links originating at n and links interfered with by transmissions from n. We consider the *link dependence graph* as defined in [2], i.e. the directed graph with vertices being links $l \in \mathcal{L}$, in which the edge from l to another vertex $l' \in \mathcal{L}$ exists if and only if $l' \in \mathcal{S}_{t_l}$. Throughout the paper we assume the *strong connectivity* of the link dependence graph, which assumes that there exists a directed path between any two vertices.

B. Saturation Throughput Region and its Properties

Suppose the network employs a slotted-Aloha-type random access protocol. Recall that each node maintains separate queues for the packets of different flows. In each time slot t, node n attempts a transmission with probability P_n, and chooses to transmit data from queue $Q_l^{(r)}$ on link l with conditional probability $p_l^{(r)}/P_n$, where $p_l^{(r)} \geq 0$ is defined for each pair (r, l) such that $n = t_l$ and $l \in \mathcal{L}_r$. Thus, $p_l^{(r)}$ is the resulting probability of transmission of class r packets on link l, and

$$P_n = \sum_{l:n=t_l} \sum_{r:l \in \mathcal{L}_r} p_l^{(r)} \leq 1, \quad \forall\ n \in \mathcal{N}. \tag{1}$$

We define $\mathcal{P}$ to be the set of all feasible vectors of link *access probabilities* $\boldsymbol{p} = \{p_l^{(r)},\ l \in \mathcal{L}_r,\ r \in \mathcal{R}\}$. Obviously,

$$\mathcal{P} = \{\boldsymbol{p} \in [0, 1]^d : P_n \leq 1, \quad \forall\ n \in \mathcal{N}\}, \tag{2}$$

where $d = \sum_{r \in \mathcal{R}} |\mathcal{L}_r|$. Given $\boldsymbol{p} \in \mathcal{P}$, the transmission attempts by all nodes are independent, and then the resulting average successful transmission rate (or, average throughput) allocated to flow r on the link $l \in \mathcal{L}_r$ is

$$\mu_l^{(r)}(\boldsymbol{p}) = p_l^{(r)} \prod_{n \neq t_l, r_l \in \mathcal{N}_n} (1 - P_n). \tag{3}$$

We will use notation $\boldsymbol{\mu}(\boldsymbol{p}) = \{\mu_l^{(r)}(\boldsymbol{p}),\ l \in \mathcal{L}_r,\ r \in \mathcal{R}\}$.

Definition 1: We define the system *saturation throughput region* $\mathcal{M}$ as the set of all possible $\boldsymbol{\mu}(\boldsymbol{p})$, along with the vectors dominated by them, namely,

$$\mathcal{M} = \{\boldsymbol{v} \in [0,1]^d : \exists p \in \mathcal{P},\ s.t.\ \boldsymbol{v} \preceq \boldsymbol{\mu}(\boldsymbol{p})\}. \tag{4}$$

We also define the log-throughput region $\log \mathcal{M}$ by

$$\log \mathcal{M} = \{\boldsymbol{u} = \log \boldsymbol{v} : \boldsymbol{v} \in \mathcal{M}, \boldsymbol{v} \in \mathbb{R}_{++}^d\}$$

and its Pareto ("north-east") boundary as

$$[\log \mathcal{M}]^* = \{\boldsymbol{u} \in \log \mathcal{M} \colon \text{if } \boldsymbol{u} \preceq \boldsymbol{u}' \in \log \mathcal{M}, \text{then } \boldsymbol{u} = \boldsymbol{u}'\}.$$

Proposition 1: (*Follows from Lemma 1 and Theorem 2 in [2]*) The log-throughput region $\log \mathcal{M}$ is convex and the boundary $[\log \mathcal{M}]^*$ is a smooth $(d-1)$-dimensional surface in $\mathbb{R}^d$, which can be parametrized by the vectors of positive *link weights* $\boldsymbol{w} = \{w_l^{(r)},\ l \in \mathcal{L}_r,\ r \in \mathcal{R}\} \in \mathbb{R}_{++}^d$, as follows. A vector $\boldsymbol{u} \in [\log \mathcal{M}]^*$ if and only if there exists a unique (up to scaling by a positive constant) link weights vector $\boldsymbol{w} \in \mathbb{R}_{++}^d$ such that $\boldsymbol{u}$ is the unique solution of the problem

$$\max\ \boldsymbol{w} \cdot \boldsymbol{u} \quad \text{s.t. } \boldsymbol{u} \in \log \mathcal{M},$$

or an equivalent problem $\max\ \boldsymbol{w} \cdot \log \boldsymbol{v}$ s.t. $\boldsymbol{v} \in \mathcal{M}$. (Thus, the vector $\boldsymbol{w}$ is the unique outer normal vector to the region $\log \mathcal{M}$ at the boundary point $\boldsymbol{u}$.) Moreover, the unique set of access probabilities $\boldsymbol{p}$ such that $\boldsymbol{u} = \log \boldsymbol{\mu}(\boldsymbol{p})$ is given by

$$p_l^{(r)} = \frac{w_l^{(r)}}{\sum_{i \in \mathcal{S}_n} \sum_{k: i \in \mathcal{L}_k} w_i^{(k)}}, \tag{5}$$

where $n = t_l$ is the transmitter node of link l. Different vectors $\boldsymbol{u} \in [\log \mathcal{M}]^*$ have different corresponding weights vectors $\boldsymbol{w}$; this implies, in particular, that region $\log \mathcal{M}$ is *strictly* convex.

We will denote by $\boldsymbol{p}(\boldsymbol{w})$ the function given by (5), and for future reference adopt the convention that $p_l^{(r)} = 0$ when $w_l^{(r)} = 0$. This makes $\boldsymbol{p}(\boldsymbol{w})$ well defined for all $\boldsymbol{w} \in \mathbb{R}_+^d$, and not just for $\boldsymbol{w} \in \mathbb{R}_{++}^d$, because $w_l^{(r)} > 0$ guarantees that the denominator in (5) is positive as well. The important feature of expression (5) is that the denominator is essentially the sum of the weights of all links with which the transmitting node n interferes including the link originating at n itself, and so nodes can compute their access probabilities very efficiently, using limited information exchange within their local neighborhoods (see [13] and [2] for more details).

C. Queueing Dynamics

The generic queuing dynamics in the random access network described above are as follows. Here we do not discuss here how new packets arrive in the networks and how access probabilities are set, which will be specified later. Let $A^{(r)}(t)$ denote the number of exogenous data packet arrivals at the source node $l(r,1)$ of flow r in time slot t, and $Q_j^{(r)}(t)$, $j = 1, \ldots, |\mathcal{L}_r|$, be the queue length of type r packets at the transmitter node of link $l(r,j)$ at time t, where the notation $l(r,j)$ is based on the convention $Q_j^{(r)} = Q_{l(r,j)}^{(r)}$. Then,

$$Q_l^{(r)}(t+1) = \begin{cases} Q_j^{(r)}(t) + A^{(r)}(t) - h_j^{(r)}(t),\ j = 1 \\ Q_j^{(r)}(t) + h_{j-1}^{(r)}(t) - h_j^{(r)}(t),\ 1 < j < |\mathcal{L}_r| \end{cases}$$

where $h_j^{(r)} = 1$ if there is a successful transmission of a flow r packet on link $l(r,j)$ in slot t, and $h_j^{(r)} = 0$ otherwise.

IV. Dynamic Queue Back-Pressure Random Access

In this section we introduce a dynamic distributed algorithm, called *Queue Back-Pressure Random Access* (QBRA), which is the main subject of this paper. The algorithm generalizes the Queue Length Based Random Access (QRA) scheme, introduced in [13] and [2] for the special case of our model, where all routes have length one. Under QRA, nodes choose their access probabilities $\mathbf{p}$ dynamically, according to formula (5), with link weights $w_l^{(r)}$ at time t being a fixed function of the current queue length $Q_l^{(r)}(t)$. In the simplest form, $w_l^{(r)} = Q_l^{(r)}(t)$. (See [13] and [2] for more general weight functions.)

Under the QBRA algorithm, nodes also dynamically choose access probabilities $\mathbf{p}$ according to (5), with the weight $w_j^{(r)}$ of flow r on link $l(r,j)$ at time t being set to the current *queue differential*, $w_j^{(r)} = \Delta Q_j^{(r)}(t)$, defined as follows:

$$\Delta Q_j^{(r)}(t) \doteq \begin{cases} \left[Q_j^{(r)}(t) - Q_{j+1}^{(r)}(t)\right]^+, 1 \le j < |\mathcal{L}_r|, \\ Q_{|\mathcal{L}_r|}^{(r)}(t), \qquad j = |\mathcal{L}_r|. \end{cases} \tag{6}$$

As usual, we identify $\Delta Q_j^{(r)}$ and $\Delta Q_{l(r,j)}^{(r)}$, and denote by $\boldsymbol{\Delta Q}$ the vector of all $\Delta Q_j^{(r)}$ in the network.

Obviously, under QBRA a transmission of a flow r packet at time t on link $l(r,j)$ will not be attempted unless $Q_j^{(r)}(t) - Q_{j+1}^{(r)}(t) > 0$. This clearly implies that if inequality

$$Q_j^{(r)}(t) \ge Q_{j+1}^{(r)}(t) - 1, \tag{7}$$

holds for flow r on link $l(r,j)$ at time $t = 0$, it then holds for all t. In all cases considered throughout this paper, (7) in fact holds for all flows and links at time 0 and then for all t.

V. Utility based End-to-End Throughput Allocation

In this section we study the scenario in which the sources of all data flows are "saturated", i.e. they have infinite amounts of data to send. Informally, the problem is to allocate throughputs $x^{(r)}$ to flows r along their respective routes in the network by setting access probabilities of all nodes in a way that maximizes the *weighted proportional fairness* objective $\sum_r \theta^{(r)} \log x^{(r)}$, where $\theta^{(r)} > 0$ are fixed weights.

This problem was considered in [15], where two distributed iterative algorithms for setting access probabilities were proposed and proved to be optimal; these approaches and results were generalized in [5]. However, the solution approaches in [15] and [5], based on the dual and the primal algorithms in convex optimization, both need end-to-end feedback information to update variables maintained by the nodes. This may induce increased delays due to the end-to-end signaling along the route, especially in large-scale networks. Moreover, the deterministic optimization-based algorithms of [15] and [5] are oblivious to the actual queueing dynamics in the network, which also may degrade performance metrics, including delays.

The purpose of this section is to prove that the above problem can be solved by the QBRA algorithm as well. The solution is very simple. Each flow r source maintains a constant queue length $Q_1^{(r)}$, proportional to $\theta^{(r)}$, at the flow source node. Then, as we show, the dynamics of the network queues under QBRA are such that the queue lengths "converge" to the values that induce access probabilities resulting in the optimal end-to-end throughput allocation. Since QBRA uses only local message passing between "neighboring" nodes, one can say that QBRA provides a "more distributed" solution to the problem than those in [15].

The solution provided by QBRA is *asymptotically optimal* in the following sense. Queues at the source nodes are maintained equal to $\theta^{(r)}/\eta$, where $\eta > 0$ is a small scaling parameter. This means that, roughly speaking, the parameter η "scales up" all queues in the network by a large factor $1/\eta$. The optimality is achieved when η becomes infinitesimally small. Consequently, our results concern *fluid limits* of the queue length process, which are the limits of the process under $\eta\boldsymbol{Q}(t/\eta)$ space and time scaling, as $\eta \downarrow 0$.

Finally, in this section we show that QBRA also solves a more general problem, with additional, minimum end-to-end throughput requirements, $x^{(r)} \geq \lambda^{(r)}$.

A. Problem Formulation

The problem is to operate our random access network in a way such that the average *end-to-end* flow throughputs $x^{(r)}$ maximize $\sum_r \theta^{(r)} \log x^{(r)}$, where $\theta^{(r)} > 0$ are fixed parameters, while keeping all the queues in the network stable. This in particular means that we want the values of $x^{(r)}$ to be those given (as $x^{(r)} = v_1^{(r)}$) by a solution of the following optimization problem for the average link-flow throughputs $\boldsymbol{v}$:

$$\begin{aligned} \max_{\boldsymbol{v}\in\mathcal{M}} \quad & \sum_{r\in\mathcal{R}} \theta^{(r)} \log v_1^{(r)}, \\ \text{subject to} \quad & v_{j-1}^{(r)} \leq v_j^{(r)}, \\ & j = 2,\ldots,|\mathcal{L}_r|,\ r\in\mathcal{R}. \end{aligned} \tag{8}$$

Here again we use notational convention $v_j^{(r)} = v_{l(r,j)}^{(r)}$, and we will adopt similar ones accordingly throughout the paper. Since any optimal solution to (8) must be such that $\boldsymbol{v} \succ 0$, problem (8) can be equivalently written in terms of log-throughputs $\boldsymbol{u} = \log \boldsymbol{v}$:

$$\begin{aligned} \max_{\boldsymbol{u}\in\log\mathcal{M}} \quad & \sum_{r\in\mathcal{R}} \theta^{(r)} u_1^{(r)}, \\ \text{subject to} \quad & u_{j-1}^{(r)} \leq u_j^{(r)}, \\ & j = 2,\ldots,|\mathcal{L}_r|,\ r\in\mathcal{R}. \end{aligned} \tag{9}$$

Note that, given smoothness of the boundary $[\log\mathcal{M}]^*$ (see Proposition 1), any interior point $\boldsymbol{u}$ of $\log\mathcal{M}$ is strictly dominated by some boundary point $\boldsymbol{u}^* \in [\log\mathcal{M}]^*$. (We can choose any dominating point and then move it slightly within the boundary so it strictly dominates.) This implies that any optimal solution $\boldsymbol{u}^*$ to (9) must lie on the boundary $[\log\mathcal{M}]^*$. Otherwise, we could move this point within the interior of $\log\mathcal{M}$ in a direction that improves the value of the objective, while respecting the constraints. Moreover, since the region $\log\mathcal{M}$ is strictly convex by Proposition 1, the optimal solution $\boldsymbol{u}^*$ to (9) is unique. (Non-uniqueness would imply that we could choose two optimal solutions, $\boldsymbol{u}^{*,1}$ and $\boldsymbol{u}^{*,2}$; then, the middle point $\boldsymbol{u}^* = (\boldsymbol{u}^{*,1} + \boldsymbol{u}^{*,2})/2$ has same objective value, but cannot be optimal, since it is in the interior.) Then $\boldsymbol{v}^*$ such that $\boldsymbol{u}^* = \log\boldsymbol{v}^*$ is the unique solution of (8). Further, again by Proposition 1, the unique (up to scaling) outer normal vector to the smooth boundary $[\log\mathcal{M}]^*$ at point $\boldsymbol{u}^*$ has all positive components. This implies that the optimal link throughputs allocated to each flow along its route are all equal:

$$u_1^{(r)*} = \ldots = u_{|\mathcal{L}_r|}^{(r)*},\ r\in\mathcal{R}. \tag{10}$$

Otherwise, we could improve the objective by slightly "moving" $\boldsymbol{u}^*$ within the boundary $[\log\mathcal{M}]^*$.

Now, consider the Lagrangian for the problem (9):

$$\begin{aligned} L(\boldsymbol{q},\boldsymbol{u}) &= \sum_{r\in\mathcal{R}} \left(\theta^{(r)} u_1^{(r)} - \sum_{j=2}^{|\mathcal{L}_r|} q_j^{(r)} \left(u_{j-1}^{(r)} - u_j^{(r)} \right) \right), \quad (11) \\ &= \Delta\boldsymbol{q}\cdot\boldsymbol{u}, \quad (12) \end{aligned}$$

where, by convention, $\boldsymbol{q}$ is such that $q_1^{(r)} = \theta^{(r)}$ for all r, and $\Delta\boldsymbol{q}$ is the vector with components

$$\Delta q_j^{(r)} = \begin{cases} q_j^{(r)} - q_{j+1}^{(r)}, & 1 \leq j < |\mathcal{L}_r|, \\ q_{|\mathcal{L}_r|}^{(r)}, & j = |\mathcal{L}_r|. \end{cases} \tag{13}$$

Then, for any optimal dual solution $\boldsymbol{q^*}$, we must have

$$\boldsymbol{u}^* = \arg\max_{\boldsymbol{u}\in\log\mathcal{M}} \Delta\boldsymbol{q}^*\cdot\boldsymbol{u}, \tag{14}$$

from which we conclude the following additional facts. First, $\Delta\boldsymbol{q^*} \succ 0$, that is

$$\theta^{(r)} > q_2^{(r)*} > \ldots > q_{|\mathcal{L}_r|}^{(r)*} > 0,\ r\in\mathcal{R}, \tag{15}$$

because if at least one of the components of $\Delta\boldsymbol{q^*}$ would be negative or zero, no point $\boldsymbol{u}$ of $\log\mathcal{M}$ could be an optimizer in (14). (Because $\Delta\boldsymbol{q}^*\cdot\boldsymbol{u}$ could always be increased by a small change of $\boldsymbol{u}$ within $\log\mathcal{M}$.) Second, again using Proposition 1, $\Delta\boldsymbol{q^*}$ is an outer normal to $[\log\mathcal{M}]^*$ at $\boldsymbol{u}^*$. From here, finally, we conclude that the optimal dual solution $\boldsymbol{q^*}$ is unique, and it is such that $\Delta\boldsymbol{q^*}$ is the unique vector of positive weights that produces the optimal rates $\boldsymbol{v}^*$ (namely, $\boldsymbol{v}^* = \boldsymbol{\mu}(\boldsymbol{p}(\Delta\boldsymbol{q^*}))$), and satisfying additional conditions $q_1^{(r)} = \theta^{(r)}$ for all r.

B. Application of QBRA Algorithm

QBRA can be applied to solve problem (8) as follows. First we fix a small parameter $\eta > 0$. Each flow r source maintains a constant queue length $Q_1^{(r)} = \lfloor\theta^{(r)}/\eta\rfloor$, at the flow source node, where $\lfloor\cdot\rfloor$ denotes the integer part of its argument. It is always feasible to provide this condition because the source has an infinite amount of data, and it can simply add a new packet in the queue after each successful transmission from it.

Otherwise, the QBRA in the network works exactly as defined earlier.

Without loss of generality, we can assume that at time $t = 0$, the relationships (7) hold (for example, all queue lengths on each route are 0, except for that of the first queue), and so (7) holds for all t.

We consider the *fluid limit* asymptotic regime. Namely, we look at a sequence of systems, with parameter $\eta \downarrow 0$. For each system we consider the space-time rescaled queueing process $\eta \boldsymbol{Q}(t/\eta)$ in continuous time $t \geq 0$, and then consider the process-level limit of those, as $\eta \downarrow 0$. The following fact, proved essentially the same way as the analogous result in [13], roughly speaking says that any limiting process is concentrated on the family of continuous trajectories $\boldsymbol{q}(t),\ t \geq 0$, called *fluid sample paths*, and describes their basic properties. We omit the proof here - it follows essentially from the same argument as that used for the analogous result in [13].

Proposition 2 (Fluid Limit): The sequence of rescaled processes $\eta \boldsymbol{Q}(t/\eta),\ t \geq 0$, can be constructed on a common probability space in a way such that, with probability 1, the sequence of realizations has a subsequence converging uniformly on compact sets to a Lipschitz continuous trajectory $\boldsymbol{q}(t),\ t \geq 0$, called the *fluid sample path* (FSP). The family of FSPs satisfies, in particular, the following properties. For each r,

$$\theta^{(r)}(t) \equiv q_1^{(r)}(t) \geq q_2^{(r)}(t) \geq \ldots \geq q_{|\mathcal{L}_r|}^{(r)}(t) \geq 0; \tag{16}$$

and for each r and $1 < j \leq |\mathcal{L}_r|$,

$$\frac{d}{dt} q_j^{(r)}(t) = \begin{cases} v_{j-1}^{(r)}(t) - v_j^{(r)}(t),\ q_j^{(r)} > 0, \\ [v_{j-1}^{(r)}(t) - v_j^{(r)}(t)]^+,\ q_j^{(r)} = 0, \end{cases}$$

where $\boldsymbol{v}(t)$ is such that

$$\boldsymbol{v}(t) \in \arg\max_{\boldsymbol{z} \in \mathcal{M}} \Delta \boldsymbol{q}(t) \cdot \log \boldsymbol{z}, \tag{17}$$

with the vector of queue differentials $\Delta \boldsymbol{q}(t) \succeq 0$ defined analogously to (13).

Note that the ordering property (16) is the limiting version of (7), and that the key property (17) follows from the fact that QBRA uses queue differentials as link weights to set access probabilities via (5).

We denote by $\mathcal{D}$ the set of all possible FSP states $\boldsymbol{q}(t)$, i.e. those satisfying inequalities (16), and by $\partial \mathcal{D}$ the subset of those $\boldsymbol{q} \in \mathcal{D}$ with at least one zero component $\Delta \boldsymbol{q}_l^{(r)} = 0$.

C. Asymptotic Optimality

Given the properties of optimal primal and dual solutions to problem (9), $\boldsymbol{u}^*$ and $\boldsymbol{q}^*$, respectively, it follows immediately that the stationary trajectory $\boldsymbol{q}(t) \equiv \boldsymbol{q}^*$ satisfies all the FSP properties described in Proposition 2. Moreover, analogously to the way it is done in [2] for a simpler model, it is easy to see that any stationary trajectory $\boldsymbol{q}(t) \equiv \boldsymbol{q}^{**} \notin \partial \mathcal{D}$, satisfying FSP properties in Proposition 2, must be such that $\boldsymbol{q}^{**} = \boldsymbol{q}^*$, because then $\boldsymbol{q}^{**}$ satisfies Karush-Kuhn-Tucker (KKT) conditions for problem (9). This to some degree motivates the following main result of this section.

Theorem 1: Every FSP is such that $\boldsymbol{q}(t) \to \boldsymbol{q}^*$ as $t \to \infty$ and, consequently, $\boldsymbol{v}(t) \to \boldsymbol{v}^*$. The convergence is uniform on all FSPs.

Theorem 1 basically says that, when the parameter $\eta > 0$ is small, then regardless of the initial state of the queues, the queues "converge to" and stay close to the values that result via the QBRA rule for access probability assignment in the optimal end-to-end throughput allocation. The key idea of the proof of Theorem 1 is contained in the following Lemma 1, which states that essentially the Lagrangian (11) of the convex optimization problem (9) can serve as a Lyapunov function to prove the convergences and we will discuss the technical ideas below.

Lemma 1: For any FSP at any time t such that $\boldsymbol{q}(t) \in \mathcal{D} \setminus \partial \mathcal{D}$, the following holds. The value of $\boldsymbol{v}(t)$, and then $\boldsymbol{u}(t) = \log \boldsymbol{v}(t)$, is defined by (17) uniquely, and moreover

$$\boldsymbol{u}(t) = \arg\max_{\boldsymbol{u} \in \log \mathcal{M}} \Delta \boldsymbol{q}(t) \cdot \boldsymbol{u}$$

and $\boldsymbol{u}(t) \in [\log \mathcal{M}]^*$. Consequently, $(\boldsymbol{q}(t), \boldsymbol{u}(t))$ is a smooth function of time in a neighborhood of t; by (12) $L(\boldsymbol{q}(t), \boldsymbol{u}(t))$ is the value of the convex dual problem to (9) at point $\boldsymbol{q}(t)$, and

$$\sum_{r \in \mathcal{R}} \theta^{(r)} u_1^{(r)*} \leq L(\boldsymbol{q}(t), \boldsymbol{u}(t)) \leq 0; \tag{18}$$

function $L(\boldsymbol{q}, \boldsymbol{u})$ is smooth in a neighborhood of $(\boldsymbol{q}(t), \boldsymbol{u}(t))$, and has zero partial gradient on primal variables $\boldsymbol{u}$ at $(\boldsymbol{q}(t), \boldsymbol{u}(t))$:

$$\nabla_{\boldsymbol{u}} L(\boldsymbol{q}(t), \boldsymbol{u}(t)) = 0. \tag{19}$$

Finally,

$$\frac{d}{dt} L(\boldsymbol{q}(t), \boldsymbol{u}(t)) \tag{20}$$

$$= - \sum_{r \in \mathcal{R}} \sum_{j=2}^{|\mathcal{L}_r|} \left(v_{j-1}^{(r)}(t) - v_j^{(r)}(t) \right) \left(u_{j-1}^{(r)}(t) - u_j^{(r)}(t) \right) \leq 0, \tag{21}$$

and the inequality is strict unless $\boldsymbol{q}(t) = \boldsymbol{q}^*$.

Proof: If $\boldsymbol{q}(t) \in \mathcal{D} \setminus \partial \mathcal{D}$, then, by Proposition 1, in the neighborhood of this point the dependence of $\boldsymbol{v}(t)$ on $\boldsymbol{q}(t)$ is given by the explicit smooth function $\boldsymbol{v}(\boldsymbol{p}(\boldsymbol{q}))$. Obviously, the dependence $\boldsymbol{u} = \log \boldsymbol{v}$ is smooth as well. Then, all the properties described in the lemma clearly follow, using in particular the smoothness of the boundary $[\log \mathcal{M}]^*$. Inequality (21) holds because each difference $v_{j-1}^{(r)}(t) - v_j^{(r)}(t)$ obviously has the same sign as the corresponding difference $u_{j-1}^{(r)}(t) - u_j^{(r)}(t)$; all such differences cannot be simultaneously equal to 0 unless $\boldsymbol{q}(t) = \boldsymbol{q}^*$, because otherwise a stationary trajectory "sitting" at a point different from $\boldsymbol{q}^*$ would exist. ■

In addition to the key Lemma 1, we need some auxiliary results to prove Theorem 1.

Lemma 2: For any FSP and any time $t \geq 0$, there exists an arbitrarily close to t time $s > t$, such that $\Delta \boldsymbol{q}(s) \succ 0$, i.e. $\boldsymbol{q}(s) \in \mathcal{D} \setminus \partial \mathcal{D}$.

Proof: Let us call any link-route pair (l, r) such that $l \in \mathcal{L}_r$, a *virtual link*. For a given FSP, let us call (l, r) a "zero" (resp. "non-zero") virtual link at time t if $\Delta q_l^{(r)} = 0$ (resp.

> 0). Suppose there are some zero virtual links at time t. If not, then the lemma statement is trivial. Since trajectory $\boldsymbol{q}(\cdot)$ is continuous, to prove the statement of the lemma it will suffice to show that there exists a time $s > t$, arbitrarily close to t, such that at least one virtual link which was zero at t becomes non-zero at s. Indeed, this implies that we can "reset" t to s, show that yet another link becomes non-zero at a time arbitrarily close to s, and so on, until all links are non-zero. Consider two cases.

Case (a): Suppose that on one of the routes r, there is a non-zero virtual link followed by a zero one; that is $\Delta q_{j-1}^{(r)}(t) > 0$ and $\Delta q_j^{(r)}(t) = 0$. This is the situation in which a transmission on the j-th link "kills" a simultaneous transmission on the $j-1$-th link. Then, it is clearly seen from (17) that $v_j^{(r)}(t) = 0$, $v_{j-1}^{(r)}(t) > 0$, and both these functions are continuous in time at t. Then, by (17), $q_j^{(r)}$ has positive, bounded away from 0, derivative in the interval $(t, t+\epsilon)$, with small $\epsilon > 0$. In the same time interval, also by (17), the derivative of $q_{j+1}^{(r)}$ is upper bounded by an arbitrarily small $\delta > 0$, if we choose sufficiently small $\epsilon > 0$. If $j = |\mathcal{L}_r|$, then $q_{j+1}^{(r)}(t) \equiv 0$ by convention. These facts mean that $\Delta q_j^{(r)}(s) > 0$ for all $s \in (t, t+\epsilon)$. We are done with case (a).

Case (b) = [NOT Case (a)]: At time t, along each route r, there is a (possibly empty) sequence of zero virtual links at the beginning, followed by the (definitely non-empty) sequence of non-zero virtual links until the end of the route. In this case, there is at least one zero virtual link, let it be the j-th link on route r, such that it either shares a link with a non-zero virtual link, or it interferes with transmissions on a non-zero virtual link. Here the latter observation uses strong connectivity of the link dependence graph. Either way, $v_j^{(r)}(t) = 0$ and it is continuous at time t. For the first non-zero virtual link on this route, say the m-th with $m > j$, $v_m^{(r)}(t) > 0$ and is continuous at t. Then, using (17) we clearly see that, in a small interval $(t, t+\epsilon)$,

$$\frac{d}{ds}[q_{j+1}^{(r)}(s) + \ldots + q_m^{(r)}(s)] < -C,$$

for some $C > 0$ independent of ϵ; and in the same interval $\frac{d}{ds} q_j^{(r)}(s) > -\delta$, where $\delta > 0$ can be made arbitrarily small by choosing small ϵ. We conclude that for any $s \in (t, t+\epsilon)$, we must have $q_j^{(r)}(s) > q_{j'}^{(r)}(s)$ for at least one j', $j+1 \le j' \le m$, and therefore one of the virtual links from the j-th to the $m-1$-th must be non-zero at time s. ∎

Lemma 3: For any FSP, $\Delta \boldsymbol{q}(t) \succ 0$ for all $t > 0$.

Proof: In view of Lemma 2, it suffices to show that if $\Delta \boldsymbol{q}(t) \succ 0$ for $t = s > 0$, then this holds for all $t \ge s$ as well. Suppose not, and τ, $s < \tau < \infty$, is the first time after s when $\boldsymbol{q}(t)$ hits set $\partial\mathcal{D}$. This means that there exists a subset of virtual links which simultaneously become zero at time τ. However, considering the values of $v_j^{(r)}(t)$ for t close to τ, and essentially repeating the argument in the proof of Lemma 2, we can show that for at least one of those links $\Delta q_j^{(r)}(t)$ must in fact be increasing for such t, and therefore cannot hit 0 at τ - a contradiction to our assumption. ∎

Proof of Theorem 1. According to Lemma 2, for any FSP at any time $t > 0$, we are in the conditions of Lemma 1. In particular, this means that the uniform bound (18) holds. Thus, to prove the uniform convergence $\boldsymbol{q}(t) \to \boldsymbol{q}^*$ it remains to show that the negative derivative $\frac{d}{dt} L(\boldsymbol{q}(t), \boldsymbol{u}(t))$, given by (20), is bounded away from zero as long as $\boldsymbol{q}(t)$ is outside of an ϵ-neighborhood of $\boldsymbol{q}^*$. This is obvious if values of $\boldsymbol{q}(t)$ are confined to a compact set, not intersecting with $\partial\mathcal{D}$. To show that it is still the case within the entire set $\mathcal{D}\backslash\partial\mathcal{D}$, it remains to observe the following. If point $\boldsymbol{q}$ approaches an arbitrary point $\boldsymbol{a} \in \partial\mathcal{D}$, the derivative $\frac{d}{dt} L$ at $\boldsymbol{q}$ approaches $-\infty$, because for at least one virtual link, the corresponding $v_{j-1}^{(r)}$ (see (20)) approaches 0 while $v_j^{(r)}$ does not, or vice versa. Here, again, we essentially repeat the argument used for Lemma 2. ∎

D. Discussion: Key Features of FSP Dynamics under QBRA

At this point we would like to highlight key features of FSP behavior under QBRA, which make it distinct from the behavior of FSPs under "conventional" back-pressure based algorithms in multi-hop networks. Consider inclusion (17), which determines the instantaneous service rates $\boldsymbol{v}(t)$ (in the fluid limit). First, under QBRA $\boldsymbol{v}(t)$ is "chosen" within the saturation throughput region $\mathcal{M}$, as opposed to the maximum possible throughput region. Second, and this is key, $\boldsymbol{v}(t)$ is chosen so that $\Delta \boldsymbol{q}(t) \cdot \log \boldsymbol{v}(t)$ is maximized, as opposed to maximizing $\Delta \boldsymbol{q}(t) \cdot \boldsymbol{v}(t)$ under conventional back-pressure algorithms. Both features are already present in paper [13] concerned, in particular, with queueing stability of the special, *single-hop* version of QBRA; and as shown in [13] these features do not "prevent" the use of standard "sum-of-queue-squares"-type Lyapunov functions (up to some adjustments) to establish stability. However, for the *multi-hop* version of QBRA considered in this paper, the second feature ($\log z$ instead of z in the right-hand side (RHS) of (17)) makes sum-of-queue-squares type Lyapunov functions inapplicable for queueing stability proofs (at least, we did not find a way to use them), which in fact was the starting point of our work. This prompted us to take a broader view which includes both utility maximization problems of this section and (as we will see in Section VI) the queueing stability problem, within a unified framework. This led us to consider the Lagrangian-type Lyapunov function used in the proof of Theorem 1, which (as the proof shows) can be used to establish FSP convergence, despite the second feature of FSP dynamics. (This Lyapunov function does, however, rely on the convexity of the *log-throughput* region $\log \mathcal{M}$.) The generalization of the utility maximization result, which we present next, is both important in its own right and (as shown in Section VI) allows the queueing stability result to "fall out" as essentially a corollary. The use of Lagrangian-type Lyapunov functions for queueing stability problems of back-pressure type algorithms, i.e. treating such problems essentially as special cases of utility maximization, is novel and (as we show) it works in cases when the traditional approach does not – this is one of the main technical contributions of this paper.

E. Generalization: Systems with Minimum Flow Rate Requirements

In practical systems, a minimum rate lower bound is often required on the end-to-end throughput to guarantee Quality of Service of the data transfers. Accordingly, Theorem 1 can be generalized to include such additional constraints. More precisely, suppose that the end-to-end throughput allocated to each flow r needs to be at least $\lambda^{(r)} \geq 0$. Formally, the more general optimization problem which we will write directly in terms of log-throughputs $\boldsymbol{u} = \log \boldsymbol{v}$ as in (9)), is

$$\begin{array}{ll} \max\limits_{\boldsymbol{u}\in\log\mathcal{M}} & \sum\limits_{r\in\mathcal{R}} \theta^{(r)} u_1^{(r)}, \\ \text{subject to} & u_{j-1}^{(r)} \leq u_j^{(r)}, \\ & j = 2,\dots,|\mathcal{L}_r|,\ r\in\mathcal{R}, \\ & \log\lambda^{(r)} \leq u_1^{(r)},\ r\in\mathcal{R}. \end{array} \tag{22}$$

We will assume the following

Feasibility condition: *Problem (22) is feasible, and moreover all its inequality constraints can be satisfied as strict inequalities.*

This condition can be interpreted as follows. Let us denote

$$\boldsymbol{\lambda}^{no} = \{\lambda_l^{(r),no},\ l\in\mathcal{L}_r,\ r\in\mathcal{R}\}, \tag{23}$$

where $\lambda_l^{(r),no} = \lambda^{(r)}$. Obviously, $\boldsymbol{\lambda}^{no}$ is the vector of minimum rates each flows needs to receive on each link, in order for the end-to-end rates to be at least $\lambda^{(r)}$. The superscript no stands for *nominal* minimum link loads. Then, using properties of the region $\mathcal{M}$, it is easy to verify that the feasiblity condition is equivalent to the following one.

Feasibility condition (an equivalent form): *Vector $\boldsymbol{\lambda}^{no}$ is strictly inside region $\mathcal{M}$, in the sense that $\boldsymbol{\lambda}^{no} \prec \boldsymbol{v}'$ for some $\boldsymbol{v}' \in \mathcal{M}$.* In other words, the saturation throughput region $\mathcal{M}$ is large enough to provide each flow r with a rate strictly greater than $\lambda^{(r)}$ on all its links.

Now, given the feasibility condition, there exists a unique optimal solution $\boldsymbol{u}^*$ such that (10) holds, and the optimal dual solution $y^{(r)*}$, $q_j^{(r)*}$, $j = 2,\dots,|\mathcal{L}_r|$, $r\in\mathcal{R}$, where $y^{(r)*}$ are the duals corresponding to the minimum throughput constraints. The generalized version of (15) is:

$$q_1^{(r)*} \equiv \theta^{(r)} + y^{(r)*} > q_2^{(r)*} > \dots > q_{|\mathcal{L}_r|}^{(r)*} > 0,\ r\in\mathcal{R}, \tag{24}$$

$\Delta\boldsymbol{q}^* \succ 0$ is defined as in (13), and, again, $\boldsymbol{u}^* = \arg\max_{\boldsymbol{u}\in\log\mathcal{M}} \Delta\boldsymbol{q}^*\cdot\boldsymbol{u}$.

To apply QBRA in this case we use a virtual queue $Y^{(r)}$, maintained by each flow r source node. “Tokens” are added to $Y^{(r)}$ at the average rate $\lambda^{(r)}$ (tokens/slot); one token is removed from it if there are any in every slot when a packet of flow r is successfully transmitted from the source node. As opposed to the previous situation, the source node uses not the constant value $\lfloor\theta^{(r)}/\eta\rfloor$ as the queue length $Q_1^{(r)}$, but rather the variable $Q_1^{(r)}(t) = \lfloor\theta^{(r)}/\eta\rfloor + Y^{(r)}(t)$. Otherwise, the QBRA in the network works exactly the same way as defined earlier in Section IV.

An FSP now contains additional component $y^{(r)}(t)$ for each r, which is a limit of $\eta Y(t/\eta)$, and it satisfies condition

$$\frac{d}{dt}y^{(r)}(t) = \begin{cases} \lambda^{(r)} - v_1^{(r)}(t),\ q_j^{(r)} > 0, \\ [\lambda^{(r)} - v_1^{(r)}(t)]^+,\ q_j^{(r)} = 0, \end{cases}$$

in addition to (17). If we denote, by convention, $q_1^{(r)}(t) \equiv \theta^{(r)} + y^{(r)}(t)$, then the key condition (17) determining $\boldsymbol{v}(t)$ still holds.

The generalization of Theorem 1 is the following.

Theorem 2: Assume the feasibility condition. Then, uniformly on all FSPs with initial states $\boldsymbol{q}(0)$ within an arbitrary fixed compact set, $\boldsymbol{q}(t)\to\boldsymbol{q}^*$ as $t\to\infty$ and, consequently, $\boldsymbol{v}(t)\to\boldsymbol{v}^*$.

Theorem 2 both generalizes and significantly strengthens a result of [2], which applies to QBRA in a system with one-hop routes and states only that *if* convergence $\boldsymbol{q}(t)\to\boldsymbol{q}^{**}$ holds then $\boldsymbol{q}^{**} = \boldsymbol{q}^*$.

Proof of Theorem 2 is carried out analogously to that of Theorem 1. We do not provide details here, but rather just the following key points. The Lagrangian in this case, which serves as a Lyapunov function in the proof, is

$$\begin{aligned} L(\boldsymbol{q},\boldsymbol{u}) &= \sum_{r\in\mathcal{R}}\theta^{(r)}u_1^{(r)} - \sum_{r\in\mathcal{R}}\sum_{j=2}^{|\mathcal{L}_r|} q_j^{(r)}\left(u_{j-1}^{(r)} - u_j^{(r)}\right) \\ &\quad - \sum_{r\in\mathcal{R}} y^{(r)}\left(\log\lambda^{(r)} - u_1^{(r)}\right) \\ &= \Delta\boldsymbol{q}\cdot\boldsymbol{u} - \sum_{r\in\mathcal{R}} y^{(r)}\log\lambda^{(r)}, \end{aligned} \tag{25}$$

where, by convention, for those flows r with $\lambda^{(r)} = 0$, we have $y^{(r)} \equiv 0$ and $y^{(r)}\log\lambda^{(r)} = 0$. For each FSP, the bounds (18) generalize as

$$\begin{aligned} &\sum_{r\in\mathcal{R}}\theta^{(r)}u_1^{(r)*} \leq L(\boldsymbol{q}(t),\boldsymbol{u}(t)) \\ &\qquad \leq L(\boldsymbol{q}(0),\boldsymbol{u}(0)) \leq -\sum_{r\in\mathcal{R}} y^{(r)}(0)\log\lambda^{(r)}. \end{aligned} \tag{26}$$

As in the proof of Theorem 1, an important intermediate step is showing that $\Delta\boldsymbol{q}(t)\succ 0$ for all $t>0$ - this is done analogously to the arguments in Lemmas 2 and 3.

VI. Stochastic Stability of a Network with Exogenous Arrivals

We now turn to another version of our model, where flow sources do *not* have an infinite supply of data to send, but rather there is a random process of exogenous arrivals to the first queue $Q_1^{(r)}$ at the flow source node. For simplicity let us assume that each such arrival process $A^{(r)}(t),\ t = 1,2,\dots$ is i.i.d. with the average rate $\lambda^{(r)} = \mathbb{E}[A^{(r)}(t)] > 0$, and all arrival processes are independent. The independent and identically distributed (i.i.d.) and independence assumptions can be greatly relaxed. Also, it is not an accident that here we use the same symbol $\lambda^{(r)}$ for the input rate as we use used for the minimum rate bound in Section V-E; the reason will become clear shortly.

Consider such a network under the QBRA random access scheme. The question is under which conditions the queueing process $\boldsymbol{Q}(t)$, $t = 0, 1, 2, \ldots$, in the network is stable. If we assume (for further simplicity) that $\mathbb{P}\{A^{(r)}(t) = 0\} > 0$ for each r, then it is clear that $\boldsymbol{Q}(t)$ is a countable state space, irreducible, aperiodic Markov chain. By stability we mean its ergodicity.

Note that, without loss of generality we can assume that the "queueing order" relations (7) hold along each route at all times.

The main result of this section is the following

Theorem 3: Suppose input rates $\lambda^{(r)} > 0$, $r \in \mathcal{R}$, satisfy the feasibility condition given in Section V-E. Then the network queueing process is stable.

This theorem generalizes to the multi-hop setting one of the stability results in [13], which apply to the single-hop system. We emphasize again that our proof, outlined below, is substantially different from that in [13], even though both use fluid limits. (See the discussion in Section V-D.) Note that, in the setting of this section, the feasibility condition in its equivalent form (see Section V-E) has a simple and intuitive meaning: the saturation throughput region is large enough to support the nominal loads $\boldsymbol{\lambda}^{no}$ imposed on the individual network links by the traffic flows.

We will use the *fluid limit technique* to establish Theorem 3 (see [13] for an application of the technique to a random-access system, and references therein to a general theory). With this technique, we look at the fluid limit, defined analogously to the way described in Section V-B. It is important to emphasize that the QBRA algorithm in the network does *not* use parameter η in any way. This is true for the use of QBRA in Section V as well, but there traffic sources use the parameter η to decide when to send new packets. In this section, the parameter η is used only to define the fluid limit asymptotic regime.

In our case, the FSPs turn out to satisfy the same properties as those for the FSPs in Section V-E, but specialized to the case $\theta^{(r)} = 0$ for all r. This is not merely coincidental - it is easy to observe that if in Section V-E we were to assume that all $\theta^{(r)} = 0$, then the behavior of each virtual queue $Y^{(r)}$ there would be analogous to the behavior of actual queue $Q_1^{(r)}$ in the setting of this section.

Then, according to fluid limit technique, to prove Theorem 3 it suffices to prove the following

Theorem 4: There exists $T > 0$ such that, uniformly on all FSPs with $\|\boldsymbol{q}(0)\| = 1$, we have $\boldsymbol{q}(t) = 0$ for all $t \geq T$.

Proof. This proof is, again, analogous to the proof of the convergence results in Theorems 1 and 2. We omit full details, but the key points are as follows. Since all $\theta^{(r)} = 0$, and consequently $y^{(r)}(t) \equiv q_1^{(r)}(t)$, the Lagrangian in (25) specializes to

$$\begin{aligned} L(\boldsymbol{q}, \boldsymbol{u}) &= -\sum_{r\in\mathcal{R}} \sum_{j=2}^{|\mathcal{L}_r|} q_j^{(r)} \left(u_{j-1}^{(r)} - u_j^{(r)}\right) \\ &\quad - \sum_{r\in\mathcal{R}} q_1^{(r)} \left(\log \lambda^{(r)} - u_1^{(r)}\right) \\ &= \Delta\boldsymbol{q} \cdot \boldsymbol{u} - \sum_{r\in\mathcal{R}} q_1^{(r)} \log \lambda^{(r)}. \end{aligned} \tag{27}$$

This Lagrangian is used as Lyapunov function, and for each

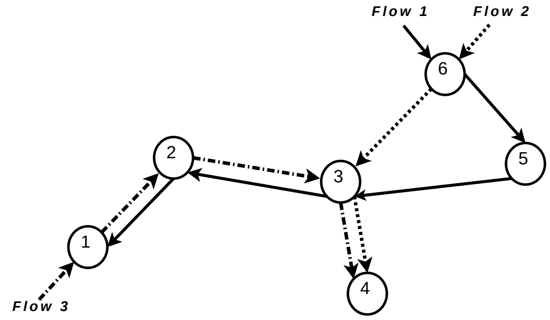

Fig. 1. A 6-node ad-hoc network

FSP we have the bounds

$$0 \leq L(\boldsymbol{q}(t), \boldsymbol{u}(t)) \leq -\sum_{r\in\mathcal{R}} q_1^{(r)}(0) \log \lambda^{(r)}. \tag{28}$$

In particular, if $\|\boldsymbol{q}(0)\| = 1$, then $L(\boldsymbol{q}(t), \boldsymbol{u}(t)) \leq -\sum_r \log \lambda^{(r)}$.

Using arguments analogous to those in Lemmas 2 and 3, we can show that for all $t > 0$ a subset of components of $\Delta\boldsymbol{q}(t)$ cannot hit 0, unless all components hit 0 simultaneously; this implies that $\Delta\boldsymbol{q}(t) \succ 0$ for all $0 < t < t'$, where t' is the first, possibly finite time when $\Delta\boldsymbol{q}(t) = 0$, and then $\boldsymbol{q}(t) = 0$. For all $0 < t < t'$, we have, analogously to (20),

$$\begin{aligned} &\frac{d}{dt} L(\boldsymbol{q}(t), \boldsymbol{u}(t)) \\ &= -\sum_{r\in\mathcal{R}} \left(\lambda^{(r)} - v_1^{(r)}(t)\right) \left(\log \lambda^{(r)} - u_1^{(r)}(t)\right) \\ &\quad - \sum_{r\in\mathcal{R}} \sum_{j=2}^{|\mathcal{L}_r|} \left(v_{j-1}^{(r)}(t) - v_j^{(r)}(t)\right) \left(u_{j-1}^{(r)}(t) - u_j^{(r)}(t)\right) \end{aligned} \tag{29}$$

Next, we prove that the RHS of (29) not only is non-positive, but in fact is bounded away from 0 by a negative constant $-\epsilon$, uniformly on all possible $\boldsymbol{u} \in [\log \mathcal{M}]^*$. Indeed, there exists a sufficiently small $\delta > 0$ such that, for any $\boldsymbol{u} \in [\log \mathcal{M}]^*$, we have

$$|\lambda^{(r)} - v_1^{(r)}| > \delta \quad \text{or} \quad |v_{j-1}^{(r)} - v_j^{(r)}| > \delta \tag{30}$$

in at least one of the terms in the RHS of (29). Otherwise $\boldsymbol{\lambda}^{no} \in [\log \mathcal{M}]^*$, which contradicts the feasibility condition. For the term that corresponds to the latter form in (30):

$$\left(v_{j-1}^{(r)} - v_j^{(r)}\right) \left(u_{j-1}^{(r)} - u_j^{(r)}\right) \geq \left(v_{j-1}^{(r)} - v_j^{(r)}\right)^2 \geq \delta^2. \tag{31}$$

Here we used the fact that $(d/dz) \log z > 1$ when $z \in (0, 1)$.

Thus, $L(\boldsymbol{q}(t), \boldsymbol{u}(t))$, and then $\boldsymbol{q}(t)$, must hit 0 within a uniformly bounded time. The fact that $\boldsymbol{q}(t)$ cannot leave 0 after first hitting it clearly follows. ■

VII. Numerical Example

In this section we investigate performance of QBRA via simulations. We consider a simple 6-node, 3-route ad hoc network as shown in Figure 1, which has the same network topology as the second example in [15]. The nodes are labelled from 1 to 6, and the interference model is such that each node interferes with the reception at its one-hop neighbor nodes in

TABLE I
THE END-TO-END THROUGHPUT $v^{(r)}$ OF EACH ROUTE IN TWO CASES: WITH/WITHOUT MINIMUM RATE REQUIREMENTS

Thru	$v^{(1)}$	$v^{(2)}$	$v^{(3)}$
Case 1	0.05196	0.12258	0.08770
Case 2	0.09934	0.07392	0.04957

the network graph. Therefore, we have

$$\mathcal{D}_1 = \{2\}, \mathcal{D}_2 = \{1, 3\}, \mathcal{D}_3 = \{2, 4\},$$
$$\mathcal{D}_4 = \emptyset, \mathcal{D}_5 = \{3\}, \mathcal{D}_6 = \{3, 5\},$$
$$\mathcal{N}_1 = \{1, 2\}, \mathcal{N}_2 = \{1, 2, 3\}, \mathcal{N}_3 = \{2, 3, 4, 5, 6\},$$
$$\mathcal{N}_4 = \{3, 4\}, \mathcal{N}_5 = \{3, 5, 6\}, \mathcal{N}_6 = \{3, 5, 6\},$$

and the sets $\mathcal{S}_n$ defined accordingly. Links are identified by the pair (r, j), so that, for example, $Q_2^{(1)}$ is the queue length of flow 1 at the second link (from node 5 to node 3) on its route.

We apply QBRA to solve both the optimal end-to-end flow throughput allocation problem of Section V, and to provide stability of the queues in the case of exogenous arrivals (Section VI). We compare the performance of QBRA with that of deterministic optimization based algorithms referred to as OPTs, such as that in [15].

We would like to emphasize that, even when QBRA and OPTs are applied to solve the same problem, such as (9), there are significant differences between them: QBRA updates network variables based on queue-lengths, while OPTs are oblivious of the actual queues; QBRA can be implemented by nodes exchanging queueing information within local neighborhoods, while OPTs require end-to-end message passing along each flow route. When we talk about providing queueing stability in a system with exogenous arrivals, the difference is even more pronounced: OPTs would require estimation of the flow input rates to be used in the appropriate optimization problem to calculate link access probabilities resulting in sufficient link throughputs along each route; QBRA does not need to know or estimate input rates and ensures stability "automatically" when feasible.

A. End-to-end Throughput Allocation

Here we run QBRA to solve the problem (22) with weights $\theta^{(1)} = \theta^{(2)} = \theta^{(3)} = 1$, in two cases. Case 1 is without minimum rate constraints; Case 2 is with minimum rate constraint $\lambda^{(1)} = 0.1$ for flow 1, and none for the other two flows. The resulting end-to-end throughputs after queues "converge" are given in Table I. These throughputs are the same as those produced by OPT as expected, since the problem being solved by the two algorithms here is same.

For both Cases 1 and 2, we run the algorithm with three different values of the scaling parameter η, namely, 0.002, 0.01 and 0.05, to demonstrate how the dynamics of queues depends on this parameter. Figures 2 and 3 show the dynamics of the queues of flow 1 along its route. As predicted by the asymptotic results in Theorems 1 and 2, the scaled queue lengths $\eta Q_j^{(1)}$ "converge" and "stabilize" around the corresponding values $q_j^{(1)*}$. Here "stabilize" and "converge" are in

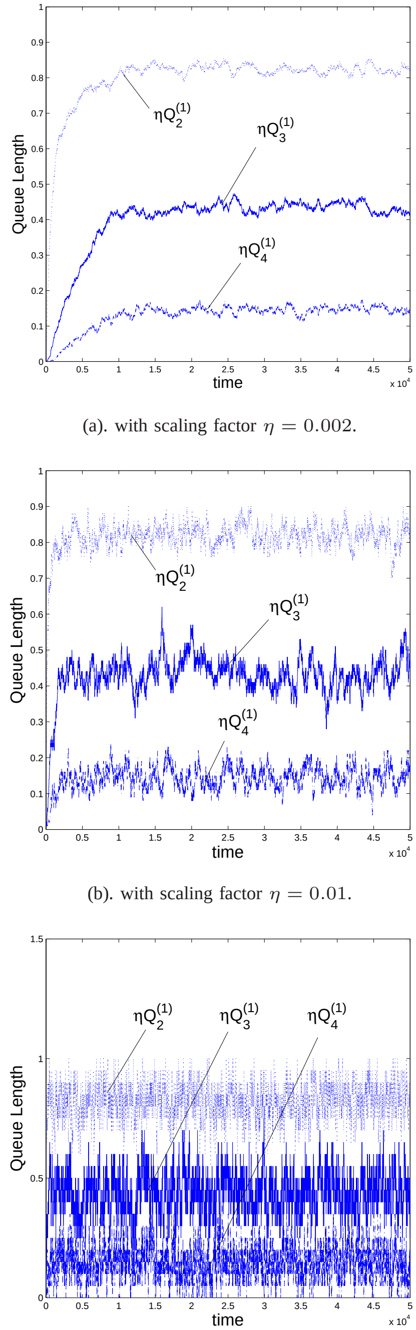

(a). with scaling factor $\eta = 0.002$.

(b). with scaling factor $\eta = 0.01$.

(c). with scaling factor $\eta = 0.05$.

Fig. 2. Behavior of flow 1 queues, under different scaling factor η.

quotation marks, because for any pre-limit system, with finite η, there cannot be a convergence in the deterministic sense - the queueing processes remain random. We can also see, again as predicted by Theorems 1 and 2, that the "convergence" time of the QBRA algorithm is roughly proportional to $1/\eta$. This is true, however, as long as η remains sufficiently small; if η is "too large" as is $\eta = 0.05$ on Figures 2(c) and 3(c), the fluctuations of the scaled queue lengths $\eta Q_j^{(r)}$ around $q_j^{(r)*}$, even after the queues "converge", will be too large, and the accuracy of the algorithm will suffer. Therefore, the value of parameter η has to be chosen carefully to achieve a balanced

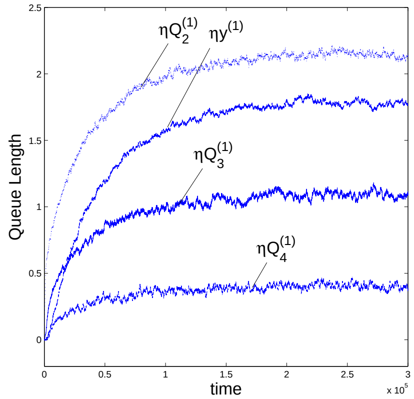

(a). with scaling factor $\eta = 0.002$.

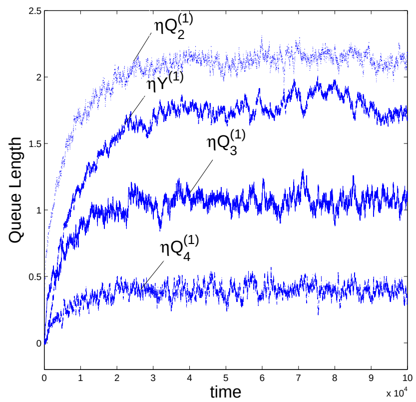

(b). with scaling factor $\eta = 0.01$.

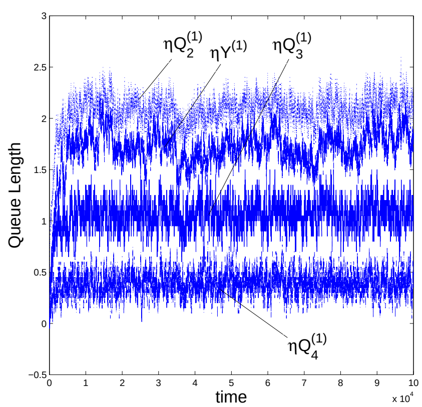

(c). with scaling factor $\eta = 0.05$.

Fig. 3. Behavior of flow 1 queues, under different scaling factor η, with minimum rate requirements $\boldsymbol{\lambda} = (0.1, 0, 0)$.

tradeoff between oscillation around stationary regime and the time to converge. Namely, it should be chosen "as large as possible, but not larger."

B. System with Exogenous Arrivals

Here we simulate the system with exogenous (i.i.d. Poisson) arrivals with equal rates for all flows, $\lambda^{(1)} = \lambda^{(2)} = \lambda^{(3)}$, and we scale the rates up to observe the changes of the queue lengths. The QBRA works exactly as specified in Section VI. An OPT algorithm that we simulate works as follows: we a priori "pre-calculate" link access probabilities so that the resulting end-to-end rates $x^{(r)}$ provided to the flows are maximal, subject to $x^{(1)} = x^{(2)} = x^{(3)}$. In other words, we pretend that an optimization based algorithm is run a priori to calculate appropriate access probabilities. OPT is oblivious to the queue lengths, except that, when there are no packets at a link, the link does *not* attempt transmission. We study the total average queue length of each flow r which by Little's law is proportional to the end-to-end queueing delay: $Q^{(r)} = \sum_l Q_l^{(r)}$. Figure 4 compares $Q^{(1)}, Q^{(2)}$ and $Q^{(3)}$ under the QBRA and OPT. It shows that the average queues under QBRA are significantly lower than those under OPT. An intuitive explanation of this is that QBRA "better adapts" to the current queue length in the network.

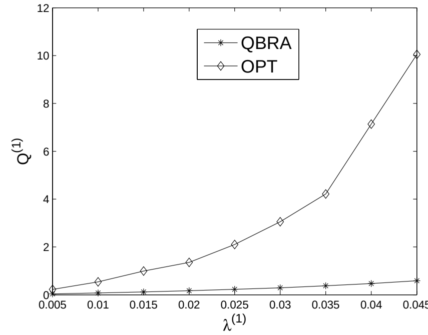

(a). Route 1: $Q^{(1)} = Q_1^{(1)} + Q_2^{(1)} + Q_3^{(1)} + Q_4^{(1)}$.

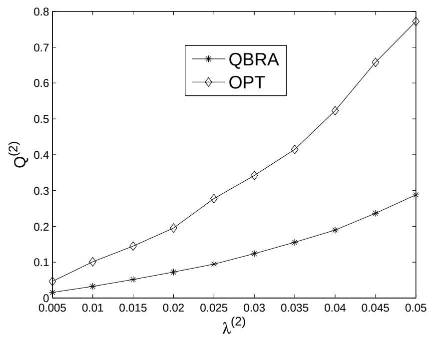

(b). Route 2: $Q^{(2)} = Q_1^{(2)} + Q_2^{(2)}$.

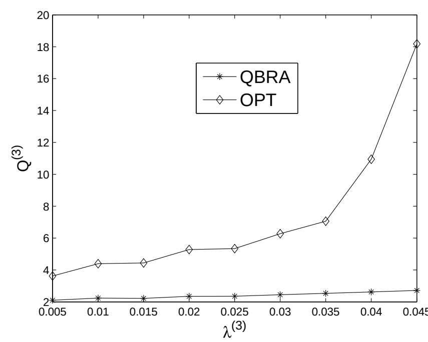

(c). Route 3: $Q^{(3)} = Q_1^{(3)} + Q_2^{(3)} + Q_3^{(3)}$.

Fig. 4. Comparison of the QBRA scheme and the optimization-based scheme on queueing performance: $\lambda^{(1)} = \lambda^{(2)} = \lambda^{(3)}$.

VIII. Conclusion and Future Work

In this paper we have considered a class of queue back-pressure random access (QBRA) algorithms within a model of wireless networks with multi-hop flow routes, where the actual queue lengths of the flows in each node's close neighborhood are used to determine the nodes' channel access probabilities. We have investigated the properties and performance of the QBRA scheme under two different traffic models.

For the model with infinite backlog at each flow source, we have shown that QBRA, combined with simple congestion control local to each source, leads to the optimal solution of a utility-based end-to-end throughput allocation, within the network *saturation throughput region* achievable by random access. The implementation of this scheme needs no end-to-end message passing as in contrast to existing pure optimization-based algorithms. We have further generalized this local QBRA scheme to the case of additional, minimum flow rate constraints. On the other hand, for the model with stochastic exogenous arrivals, we have shown that QBRA ensures stochastic stability of the queueing process as long as nominal loads of the nodes are within the saturation throughput region.

One subject of interest for future work is a study of the queueing stability of random access schemes in the model of multihop transmissions with link weights being defined more generally than queue differentials. Meanwhile, another topic is of interest to quantify and compare the queue performance under different queue based random access schemes, and thus to determine the optimal queue function in terms of queueing delay.

References

[1] N. Abramson. The ALOHA system - another alternative for computer communications. In *Proceedings of AFIPS Conference*, volume 37, pages 281–285, Fall Joint Comupter Conference, Montvale, NJ, 1970.

[2] P. Gupta and A. Stolyar. Optimal throughput allocation in general random-access networks. In *Proceedings of the 38th Conference on Information Sciences and Systems*, Princeton, NJ, March 2006.

[3] IEEE. Wireless LAN medium access control (MAC) and physical layer (PHY) specifications IEEE standard 802.11., June 2006. IEEE Standard 802.11.

[4] C. Joo and N. B. Shroff. Performance of random access scheduling schemes in multi-hop wireless networks. In *Proceedings of IEEE Infocom*, Anchorage, Alaska, 2007.

[5] J. W. Lee, M. Chiang, and A. R. Calderbank. Jointly optimal congestion and contention control in wireless ad hoc networks. *IEEE Communications Letters*, 10(3):216–218, 2006.

[6] J. W. Lee, M. Chiang, and A. R. Calderbank. Utility-optimal random-access control. *IEEE Transactions on Wireless Communications*, 6(7):2741–2751, July 2007.

[7] X. Lin and S. Rasool. Constant-time distributed scheduling policies for ad hoc wireless networks. In *Proceedings of IEEE Conference on Decision and Control*, San Diego, CA, 2006.

[8] E. Modiano, D. Shah, and G. Zussman. Maximizing throughput in wireless network via gossiping. In *Proceedings of ACM Sigmetrics*, Saint-Malo, France, 2006.

[9] A. H. Mohsenian-Rad, J. Huang, M. Chiang, and V. Wong. Utility-optimal random access: Optimal performance without frequent explicit message passing, 2007. Submitted to *IEEE Transactions on Wireless Communications*.

[10] A. H. Mohsenian-Rad, J. Huang, M. Chiang, and V. Wong. Utility-optimal random access: Reduced complexity, fast convergence, and robust performance, 2007. Submitted to *IEEE Transactions on Wireless Communications*.

[11] T. Nandagopal, T.-E. Kim, X. Gao, and V. Bharghavan. Achieving MAC layer fairness in wireless packet networks. In *Proceedings of ACM MobiCom*, pages 87–98, Boston, MA, 2000.

[12] S. Sanghavi, L. Bui, and R. Srikant. Distributed link scheduling with constant overhead. In *Proceedings of ACM Sigmetrics*, San Diego, CA, 2007.

[13] A. Stolyar. Dynamic distributed scheduling in random access network. *Journal of Applied Probability*, 45(2):297–313, June 2008.

[14] L. Tassiulas and A. Ephremides. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 37:1936–1948, 1992.

[15] X. Wang and K. Kar. Cross-layer rate control for end-to-end proportional fairness in wireless networks with random access. In *Proceedings of MobiHoc*, Urbana-Champaign, IL, 2005.

[16] Y. Yi, A. Proutière, and M. Chiang. Throughput of random access without message passing. In *Proceedings of the 42th Conference on Information Sciences and Systems*, Princeton, NJ, 2008.

PLACE PHOTO HERE

Jiaping Liu received her bachelor's degree from Shanghai Jiao Tong University in 2004, M.A. from Princeton University in 2006, Ph.D. from Princeton University in 2009, all in Electrical Engineering. Her Ph.D. thesis focused on stochastic modeling and analysis of data networks, and its applications in wireless ad hoc networks, in particular the random access algorithm design of medium access control protocols. She was a research assistant at Bell Laboratories in summer 2007 and visiting scholar at Department of Electrical Engineering at Stanford University in fall 2007. Her Ph.D. work was partially supported by the Gordon Wu Fellowship of Princeton University.

PLACE PHOTO HERE

Alexander Stolyar is a Distinguished Member of Technical Staff in the Industrial Mathematics and Operations Research Department of Bell Labs, Alcatel-Lucent, Murray Hill, New Jersey. He received Ph.D. in Mathematics from the Institute of Control Sciences, USSR Academy of Science, Moscow, in 1989. Before joining Bell Labs in 1998, he was with the Institute of Control Sciences (Moscow), Motorola (Arlington Heights, IL) and AT&T Labs-Research (Murray Hill, NJ). His research interests are in stochastic processes, queueing theory, and stochastic modeling of communication, especially wireless, systems. He developed models and algorithms for resource allocation and QoS support in wireless systems, and holds several patents. He is a member of INFORMS; in 2004 he received the Bi-Annual Best Publication Award from the INFORMS Applied Probability Society.

PLACE PHOTO HERE

Mung Chiang (S00-M03-SM08) is an Associate Professor of Electrical Engineering, and an Affiliated Faculty of Applied and Computational Mathematics and of Computer Science, at Princeton University. He received the B.S. (Honors) in Electrical Engineering and Mathematics, M.S. and Ph.D. degrees in Electrical Engineering from Stanford University in 1999, 2000, and 2003, respectively, and was an Assistant Professor at Princeton University 2003-2008. His research areas include optimization, distributed control, and stochastic analysis of communication networks, with applications to the Internet, wireless networks, broadband access networks, and content distribution.

His awards include Presidential Early Career Award for Scientists and Engineers 2008 from the White House, Young Investigator Award 2007 from ONR, TR35 Young Innovator Award 2007 from Technology Review, Young Researcher Award Runner-up 2004-2007 from Mathematical Programming Society, CAREER Award 2005 from NSF, as well as Frontiers of Engineering Symposium participant 2008 from NAE and SEAS Teaching Commendation 2007 from Princeton University. He was a Princeton University Howard B. Wentz Junior Faculty and a Hertz Foundation Fellow. His paper awards include ISI citation Fast Breaking Paper in Computer Science, IEEE INFOCOM Best Paper Finalist, and IEEE GLOBECOM Best Student Paper. His guest and associate editorial services include IEEE/ACM Trans. Netw., IEEE Trans. Inform. Theory, IEEE J. Sel. Area Comm., IEEE Trans. Comm., IEEE Trans. Wireless Comm., and J. Optimization and Engineering, and he co-chaired 38th Conference on Information Sciences and Systems.

PLACE PHOTO HERE

H. Vincent Poor (S72, M77, SM82, F87) received the Ph.D. degree in EECS from Princeton University in 1977. From 1977 until 1990, he was on the faculty of the University of Illinois at Urbana-Champaign. Since 1990 he has been on the faculty at Princeton, where he is the Dean of Engineering and Applied Science, and the Michael Henry Strater University Professor of Electrical Engineering. Dr. Poor's research interests are in the areas of stochastic analysis, statistical signal processing and their applications in wireless networks and related fields. Among his publications in these areas are the recent books *MIMO Wireless Communications* (Cambridge University Press, 2007), co-authored with Ezio Biglieri, et al, and *Quickest Detection* (Cambridge University Press, 2009), co-authored with Olympia Hadjiliadis.

Dr. Poor is a member of the National Academy of Engineering, a Fellow of the American Academy of Arts and Sciences, and a former Guggenheim Fellow. He is also a Fellow of the Institute of Mathematical Statistics, the Optical Society of America, and other organizations. In 1990, he served as President of the IEEE Information Theory Society, and in 2004-07 as the Editor-in-Chief of these Transactions. He is the recipient of the 2005 IEEE Education Medal. Recent recognition of his work includes the 2007 IEEE Marconi Prize Paper Award, the 2007 Technical Achievement Award of the IEEE Signal Processing Society, and the 2008 Aaron D. Wyner Distinguished Service Award of the IEEE Information Theory Society.