Distributed Queue-Length based Algorithms for Optimal End-to-End Throughput Allocation and Stability in Multi-hop Random Access Networks

Jiaping Liu
Department of Electrical Engineering
Princeton University
Princeton, NJ 08544, USA
jiapingl@princeton.edu

Alexander L. Stolyar
Bell Labs, Alcatel-Lucent
600 Mountain Ave, 2C-322
Murray Hill, NJ 07974, USA
stolyar@research.bell-labs.com

***Abstract*— We consider a model of wireless network with random (slotted-Aloha-type) access and with multi-hop flow routes. The goal is to devise distributed strategies for optimal utility-based end-to-end throughput allocation and queueing stability. We consider a class of *queue back-pressure random-access algorithms* (QBRA), where actual queue lengths of the flows (in each node's close neighborhood) are used to determine nodes' channel access probabilities. This is in contrast to previously proposed algorithms, which are purely optimization-based and oblivious of actual queues. QBRA is also substantially different from much studied "MaxWeight" type scheduling algorithms, also using back-pressure.**

For the model with infinite backlog at each flow source, we show 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 a random access. (No end-to-end message passing is required.) This scheme generalizes for the case of additional, minimum flow rate constraints. For the model with stochastic exogenous arrivals, we show 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 (queue oblivious) optimization-based random access algorithms, shows that QBRA performs better in terms of end-to-end delays.

I. Introduction

In wireless ad hoc networks, contention resolution and interference avoidance 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 and BEB mechanism for contention control can be inefficient (eg. [6]). 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 (message passing) overhead.

It has been shown that the maximum throughput region can be achieved by much studied "MaxWeight"-type scheduling algorithms (originally proposed in [8]). However, in the context of wireless networks, MaxWeight algorithms typically need to be centralized. 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 with multi-hop transmissions.

Random access models have been widely adopted in contemporary works [1], [2], [4], [5], [7], [9]. Informally, we can classify them into two categories: "pure optimization-based" algorithms (eg. [4], [5], [9]) and dynamic, queue-length based strategies (eg. [2], [7]). Algorithms of the former type (eg. [4], [5], [9]) 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 a priori specified and are typically oblivious of the dynamics of actual queues in the network. The latter type algorithms ([2], [7]), even in cases when they have same optimization objective as the former algorithms, do the optimization by adaptively responding to actual queueing dynamics; for example, they do not need a priori knowledge of traffic flow input rates to achieve queueing stability (if such is feasible).

In this paper we propose and study a class of *queue back-pressure random access* (QBRA) algorithms for a multi-hop network. The algorithms use flow queue differentials on the links to determine link access probabilities. (The MaxWeight algorithms use queue differentials as well, but in a completely different way.) Our main contributions are as follows.

(i) For the problem of utility-based end-to-end throughput allocation (in Section V), when traffic sources have infinite backlog of data to send, we prove that QBRA combined with extremely simple congestion controller 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 much strengthens the corresponding result of [2], which applies to single-hop flows. A further generalization - to more general utility functions - is also possible; this is subject of future work.)

(ii) For the problem of queueing stability in a system with stochastic exogenous arrivals (in Section VI), we prove

This research was in part supported by the DARPA CBMANET project

that QBRA "automatically" (without knowing input rates) ensures stability, as long as nominal link loads are within the network *saturation throughput region*. (This generalizes some of the stability results of [7], which apply to single-hop flows.)

(iii) Finally, we present simulation results (in Section VII) showing good performance of QBRA in terms of end-to-end delays.

II. Basic Notation and Definitions

Typically, we use bold letters $\boldsymbol{x}, \boldsymbol{y}, \dots$ to denote vectors, as opposed to scalars x, y. 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 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 of link l, respectively. 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. (We will use term flow and route interchangeably, and use index r 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 keeps 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 j-th node in its route.

The system operates in discrete (or, slotted) time $t = 0, 1, 2, \dots$. In any time slot, each node may attempt to transmit one packet on (at most) one of its outgoing links. A packet transmission attempt on a link is successful, if it is not "interfered by" another simultaneous (same slot) transmission; otherwise the transmission fails. The interference model is same as in [2], [7] (and is somewhat more general than in [9]). 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 $\mathcal{N}_n \subseteq \mathcal{N}$ it interferes with, 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 such that 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 by transmissions from n.) We consider the *link dependence graph* as defined in [2], i.e. the directed graph with vertices being links $l \in \mathcal{L}$, and 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. Recall that each node keeps 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 $\mathcal{P}$ as 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 log-throughput region $\log \mathcal{M}$ as

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

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

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

Proposition 1 ([2]): Log-throughput region $\log \mathcal{M}$ is strictly convex and the boundary $[\log \mathcal{M}]^*$ is a smooth $(d-1)$-dimensional surface in $\mathbb{R}^d$, which can be parameterized 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. Vector $\boldsymbol{u} \in [\log \mathcal{M}]^*$ if and only if there exists 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. $\boldsymbol{v} \in \mathcal{M}$. 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.

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 $\boldsymbol{w} \in \mathbb{R}_{++}^d$, 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 the transmitting node n interferes with (plus the link originating at n itself), and so nodes can compute their access probabilities very efficiently, using limited information exchange within their local neighborhood. (See [2], [7] for more details.)

C. Queueing Dynamics

The generic queuing dynamics in the random access network described above is as follows. (We do not discuss here how new packets arrive in the networks and how access probabilities are set. This 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 $Q_j^{(r)}(t)$, $j = 1, \ldots, |\mathcal{L}_r|$, is the queue length of type r packets at the (transmitter node of) link $l(r,j)$ at time t. (Recall 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-Backpressure Random Access

In this section we introduce a dynamic distributed algorithm, called *Queue-Backpressure Random Access* (QBRA), which is the main subject of this work. The algorithm generalizes Queue Length Based Random Access (QRA) scheme, introduced in [2], [7] 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 its current queue length $Q_l^{(r)}(t)$. In the simplest form, $w_l^{(r)} = Q_l^{(r)}(t)$. (See [2], [7] for more general weight functions.)

Under 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*, defined as follows:

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

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 easily 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 where the sources of all data flows are "saturated", i.e. have infinite amount 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 *weighted proportional fairness* objective $\sum_r \theta^{(r)} \log x^{(r)}$, where $\theta^{(r)} > 0$ are fixed weights.

This problem was considered in [9], where two distributed iterative algorithms for setting access probabilities were proposed and proved optimal; these approaches and results were generalized in [4]. However, the solution approaches in [4], [9], based on the 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 optimization-based algorithms of [4], [9] are oblivious of 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 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 is such that the queue length "converge" to the values that induce access probabilities resulting in the optimal end-to-end throughput allocation. Since QBRA only uses local message passing between "neighboring" nodes, one can say that QBRA provides a "more distributed" solution to the problem, than those in [9].

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, parameter η "scales up" all queues in the network by large factor $1/\eta$. The optimality is achieved when η becomes infinitely small. Consequently, our results concern *fluid limits* of the queue length process, which are (roughly) 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, $x^{(r)} \ge \lambda^{(r)}$.

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 $x^{(r)}$ to be those given (as $x^{(r)} = v_1^{(r)}$) by a solution of the following optimization problem for the average link-flow throughputs $\boldsymbol{v}$:

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

(Here we use notational convention $v_j^{(r)} = v_{l(r,j)}^{(r)}$, and we use similar ones later in 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 $\boldsymbol{u} = \log \boldsymbol{v}$:

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

Since region $\log\mathcal{M}$ is strictly convex by Proposition 1, the optimal solution $\boldsymbol{u}^*$ to (9) is unique. (And then $\boldsymbol{v}^*$ such that $\boldsymbol{u}^* = \log \boldsymbol{v}^*$ is the unique solution of (8). Moreover, again using Proposition 1 and considering the Lagrangian for the problem (9) (given in (15) below), it is easy to establish the following facts. The optimal link throughputs allocated to each flow along its route are all equal:

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

the optimal dual variables $q_j^{(r)*}$, $j = 2, \ldots, |\mathcal{L}_r|$, $r \in \mathcal{R}$, corresponding to the inequality constraints in (9), are unique and are such that:

$$\theta^{(r)} > q_2^{(r)*} > \ldots > q_{|\mathcal{L}_r|}^{(r)*} > 0, \ r \in \mathcal{R}. \tag{11}$$

We will put by convention $q_1^{(r)*} = \theta^{(r)}$, and use vector notation $\boldsymbol{q}^* = \{q_l^{(r)*}, \ l \in \mathcal{L}_r, \ r \in \mathcal{R}\}$. We also denote by $\Delta\boldsymbol{q}^* \succ 0$ the vector with components

$$\Delta q_j^{(r)*} = \begin{cases} q_j^{(r)*} - q_{j+1}^{(r)*}, & 1 \le j < |\mathcal{L}_r|, \\ q_{|\mathcal{L}_r|}^{(r)*}, & j = |\mathcal{L}_r|. \end{cases} \tag{12}$$

Again considering the Lagrangian for problem (9), it is easy to see that $\boldsymbol{u}^* = \arg\max_{\boldsymbol{u}\in\log\mathcal{M}} \Delta\boldsymbol{q}^* \cdot \boldsymbol{u}$, which means that $\Delta\boldsymbol{q}$ is nothing else but the (unique) vector of weights, which results in the optimal rates $\boldsymbol{v}^*$, i.e. $\boldsymbol{v}^* = \boldsymbol{\mu}(\boldsymbol{p}(\Delta\boldsymbol{q}))$.

B. Application of QBRA algorithm. Fluid Limit

The application of QBRA to solve problem (8) is as follows. A (small) parameter $\eta > 0$ is fixed. Each flow r source maintains a constant queue length $Q_1^{(r)} = \lfloor \theta^{(r)}/\eta \rfloor$, at the flow source node. (Here $\lfloor \cdot \rfloor$ is the integer part of of a number.) It can always do that, because the source has 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 relations (7) hold (for example, all queues on each route are 0, except the first queue), and so (7) holds for all t.

We consider the *fluid limit* asymptotic regime. Namely, we look at a sequence of system, with parameter $\eta \downarrow 0$. For each system we consider the space-time rescaled queueing process $\eta\boldsymbol{Q}(t/\eta)$ in continuous time $t \ge 0$, and then consider the process-level limit of those, as $\eta \downarrow 0$. The following fact, proved essentially same way as the analogous result in [7], roughly speaking says that any limiting process is concentrated on the family of (continuous) trajectories $\boldsymbol{q}(t), \ t \ge 0$, called *fluid sample paths*, and describes their basic properties. (We omit the proof here - it follows essentially same argument as that used for analogous result in [7].)

Proposition 2 (Fluid Limit): The sequence of rescaled processes $\eta\boldsymbol{Q}(t/\eta), \ t \ge 0$, 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 $\boldsymbol{q}(t), \ t \ge 0$, called *fluid sample path* (FSP). The family of FSPs satisfies, in particular, the following properties. For each r,

$$\theta^{(r)}(t) \equiv q_1^{(r)}(t) \ge q_2^{(r)}(t) \ge \ldots \ge q_{|\mathcal{L}_r|}^{(r)}(t) \ge 0. \tag{13}$$

For each r and $1 < j \le |\mathcal{L}_r|$,

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

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

$$\boldsymbol{v}(t) \in \arg\max_{\boldsymbol{z}\in\mathcal{M}} \Delta\boldsymbol{q}(t) \cdot \log \boldsymbol{z}, \tag{14}$$

with the vector of queue differentials $\Delta\boldsymbol{q}(t) \succeq 0$ being defined analogously to (12).

Note that ordering property (13) is the limit version of (7), and that the key property (14) 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 (13), 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, it is easy to see (analogously to the way it is done in [2] for a simpler model) 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 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 FSP.

Theorem 1 basically says that, when parameter $\eta > 0$ is small, then regardless of the initial state of the queues, the queues "converge to" and stay close to the values which result (via 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 (up to some technical work we will do below) the Lagrangian of the convex optimization problem (9),

$$L(\boldsymbol{q},\boldsymbol{u}) = \sum_{r\in\mathcal{R}} \left(\theta^{(r)} u_1^{(r)} - \sum_{j=2}^{|\mathcal{L}_r|} q_j^{(r)} \left(u_{j-1}^{(r)} - u_j^{(r)} \right) \right), \tag{15}$$

$$= \Delta\boldsymbol{q}\cdot\boldsymbol{u}, \tag{16}$$

(where, by convention, $\boldsymbol{q}$ is such that $q_1^{(r)} = \theta^{(r)}$ for all r) can serve as a Lyapunov function to prove the convergence.

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 (14) 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 (16) $L(\boldsymbol{q}(t), \boldsymbol{u}(t))$ is the value of the convex problem dual to (9) at point $\boldsymbol{q}(t)$, and

$$\sum_{r\in\mathcal{R}} \theta^{(r)} u_1^{(r)*} \le L(\boldsymbol{q}(t), \boldsymbol{u}(t)) \le 0; \tag{17}$$

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 $(\boldsymbol{q}(t), \boldsymbol{u}(t))$:

$$\nabla_{\boldsymbol{u}} L(\boldsymbol{q}(t), \boldsymbol{u}(t)) = 0. \tag{18}$$

Finally,

$$\frac{d}{dt} L(\boldsymbol{q}(t), \boldsymbol{u}(t)) = -\sum_{r\in\mathcal{R}} \sum_{j=2}^{|\mathcal{L}_r|} \left(v_{j-1}^{(r)}(t) - v_j^{(r)}(t) \right) \left(u_{j-1}^{(r)}(t) - u_j^{(r)}(t) \right) \tag{19}$$

$$\le 0, \tag{20}$$

and the inequality in (20) 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 easily follow, using in particular the smoothness of the boundary $[\log\mathcal{M}]^*$. Inequality (20) holds because each difference $v_{j-1}^{(r)}(t) - v_j^{(r)}(t)$ obviously has same sign as $u_{j-1}^{(r)}(t) - u_j^{(r)}(t)$; 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 $t \ge 0$, 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 a link-route pair (l, r) such that $l \in \mathcal{R}$ 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, the lemma statement is trivial.) Since trajectory $\boldsymbol{q}(\cdot)$ is continuous, to prove statement of the lemma it will suffice to show that there exists time $s > t$, arbitrarily close to t, such that at least one virtual link which was zero at t becomes non-zero at s. Consider two cases.

Case (a): Suppose 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 where a transmission on the j-th link "kills" a simultaneous transmission on the $j-1$-th link.) Then, it is easily seen from (14) that $v_j^{(r)}(t) = 0$, $v_{j-1}^{(r)}(t) > 0$, and both these functions are continuous in time at t. Then, by (14), $q_j^{(r)}$ has positive, bounded away from 0, derivative in the interval $(t, t+\epsilon)$, with small $\epsilon > 0$. In the same time interval, also by (14), the derivative of $q_{j+1}^{(r)}$ is upper bounded by an arbitrarily small $\delta > 0$, if we choose small enough $\epsilon > 0$. (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. (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 m-th with $m > j$, $v_m^{(r)}(t) > 0$ and is continuous at t. Then, using (14) we easily 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 j-th to $m-1$-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 - a contradiction. ■

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 (17) hold. 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 (19), 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 $\mathcal{D}\setminus\partial\mathcal{D}$, 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 (19)) approaches 0 while $v_j^{(r)}$ is not, or vice versa. (Here, again, we essentially repeat the argument of Lemma 2.) ■

D. 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, additionally, 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

$$\begin{aligned} \max_{\boldsymbol{u}\in\log\mathcal{M}} \quad & \sum_{r\in\mathcal{R}} \theta^{(r)} u_1^{(r)}, \\ \text{subject to} \quad & u_{j-1}^{(r)} \leq u_j^{(r)}, \\ & j = 2, \ldots, |\mathcal{L}_r|,\ r \in \mathcal{R}, \\ & \log \lambda^{(r)} \leq u_1^{(r)},\ r \in \mathcal{R}. \end{aligned} \tag{21}$$

We will assume that problem (21) is feasible, and moreover *all inequality constraints can be satisfied as strict inequalities*. (This will be referred to as the *feasibility condition*.) Then, there exists a unique optimal solution $\boldsymbol{u}^*$ such that (10) holds, and the optimal dual solution $y^{(r)*}$, $q_j^{(r)*}$, $j = 2, \ldots, |\mathcal{L}_r|$, $r \in \mathcal{R}$, where $y^{(r)*}$ are the duals corresponding to the minimum throughput constraints. The generalized version of (11) is:

$$q_1^{(r)*} \equiv \theta^{(r)} + y^{(r)*} > q_2^{(r)*} > \ldots > q_{|\mathcal{L}_r|}^{(r)*} > 0,\ r \in \mathcal{R}, \tag{22}$$

$\Delta\boldsymbol{q}^* \succ 0$ is defined as in (12), and, again, $\boldsymbol{u}^* = \arg\max_{\boldsymbol{u}\in\log\mathcal{M}} \Delta\boldsymbol{q}^* \cdot \boldsymbol{u}$.

Application of QBRA in this case uses 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 is any) in every slot when a packet of flow r is successfully transmitted (from source node). As opposed to the previous situation, 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 same way as before.

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 (14). If we denote, by convention, $q_1^{(r)}(t) \equiv \theta^{(r)} + y^{(r)}(t)$, then the key condition (14) determining $\boldsymbol{v}(t)$ still holds.

The generalization of Theorem 1 is the following.

Theorem 2: Assume the feasibility condition. Then, uniformly on all FSP with initial states $\boldsymbol{q}(0)$ within 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 much 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, just the following key points. The Lagrangian in this case, which serves as a Lyapunov function in the proof, is:

$$\begin{aligned} L(\boldsymbol{q}, \boldsymbol{u}) &= \sum_{r\in\mathcal{R}} \theta^{(r)} u_1^{(r)} - \sum_{j=2}^{|\mathcal{L}_r|} q_j^{(r)} \left(u_{j-1}^{(r)} - u_j^{(r)}\right) \\ &\quad - \sum_{r\in\mathcal{R}} y^{(r)} \left(\log \lambda^{(r)} - u_1^{(r)}\right) \\ &= \Delta\boldsymbol{q} \cdot \boldsymbol{u} - \sum_{r\in\mathcal{R}} y^{(r)} \log \lambda^{(r)}, \end{aligned} \tag{23}$$

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 (17) generalize as

$$\begin{aligned} \sum_{r\in\mathcal{R}} \theta^{(r)} u_1^{(r)*} &\leq L(\boldsymbol{q}(t), \boldsymbol{u}(t)) \\ &\leq L(\boldsymbol{q}(0), \boldsymbol{u}(0)) \leq -\sum_{r\in\mathcal{R}} y^{(r)}(0) \log \lambda^{(r)}. \end{aligned} \tag{24}$$

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 a 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 a the first queue $Q_1^{(r)}$ at the flow source node. For simplicity let us assume that each such arrival process $A^{(r)}(t),\ t = 1, 2, \ldots$ is i.i.d. with the average rate $\lambda^{(r)} = \mathbb{E}[A^{(r)}(t)] > 0$, and all arrival processes are independent. (The 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-D; the reason will become clear shortly.)

Consider such a network under 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 understand 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, as given in Section V-D. Then the network queueing process is stable.

This theorem generalizes to the multi-hop setting one of the stability results in [7], which apply to the single-hop system. (It should be noted that our proof, outlined below, is substantially different from that in [7], even though both use fluid limits.)

We will use the *fluid limit technique* to establish Theorem 3. (See [7] 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 parameter η to decide when to send new packets. In this section, parameter η is only used 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-D, but specialized to the case $\theta^{(r)} = 0$ for all r. (This is not an accident - it is easy to observe that if in Section V-D 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 this section setting.)

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 of Theorem 4 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 (23) specializes to

$$L(\boldsymbol{q}, \boldsymbol{u}) = -\sum_{j=2}^{|\mathcal{L}_r|} q_j^{(r)} \left(u_{j-1}^{(r)} - u_j^{(r)} \right)$$
$$- \sum_{r \in \mathcal{R}} q_1^{(r)} \left(\log \lambda^{(r)} - u_1^{(r)} \right)$$

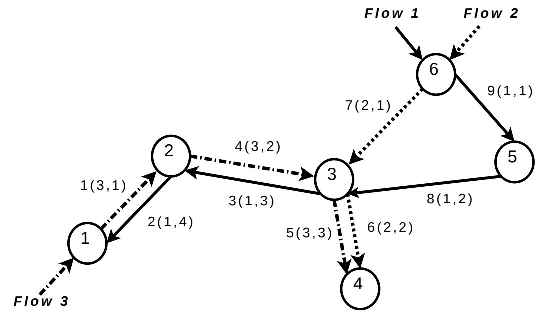

Fig. 1. A 6-node ad-hoc network

$$= \Delta \boldsymbol{q} \cdot \boldsymbol{u} - \sum_{r \in \mathcal{R}} q_1^{(r)} \log \lambda^{(r)}. \qquad (25)$$

This Lagrangian is used as Lyapunov function, and for each FSP we have the bounds

$$0 \leq L(\boldsymbol{q}(t), \boldsymbol{u}(t)) \leq -\sum_{r \in \mathcal{R}} q_1^{(r)}(0) \log \lambda^{(r)}. \qquad (26)$$

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

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 (19),

$$\frac{d}{dt} L(\boldsymbol{q}(t), \boldsymbol{u}(t))$$
$$= -\sum_{r \in \mathcal{R}} \left(\lambda^{(r)} - v_1^{(r)}(t) \right) \left(\log \lambda^{(r)} - u_1^{(r)}(t) \right)$$
$$- \sum_{r \in \mathcal{R}} \sum_{j=2}^{|\mathcal{L}_r|} \left(v_{j-1}^{(r)}(t) - v_j^{(r)}(t) \right) \left(u_{j-1}^{(r)}(t) - u_j^{(r)}(t) \right) \qquad (27)$$

Finally, we observe that the RHS of (27) not only is non-positive, but in fact bounded away from 0 by a negative constant $-\epsilon$, uniformly on all possible $\boldsymbol{u} \in [\log \mathcal{M}]^*$. 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 easily follows.

VII. Numerical Example

In this section we investigate performance of QBRA via a numerical example. 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 [9]. The nodes are labelled from 1 to 6, and the links are labelled from 1 to 9 and the route-link pair as in the format $l(r, j)$, i.e. $5(3, 3)$ denotes the link 5 in $\mathcal{L}$ which is the 3rd link of route 3. The interference model is that each node interferes with the reception at its 1-hop neighbors; for example, $\mathcal{N}_2 = \{1, 2, 3\}$, $\mathcal{N}_3 = \{2, 3, 4, 5, 6\}$.

The behavior of QBRA for the optimal end-to-end flow throughput allocation, with and without minimum throughput constraints is as predicted by Theorems 1 and 2 in Section V. We do not present those simulations to save space. Instead

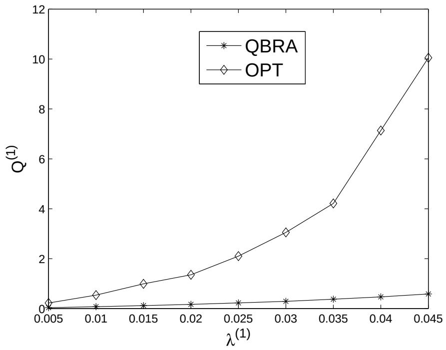

(a). Route 1: $Q^{(1)} = Q_1^{(1)} + Q_2^{(1)} + Q_3^{(1)} + Q_4^{(1)}$.

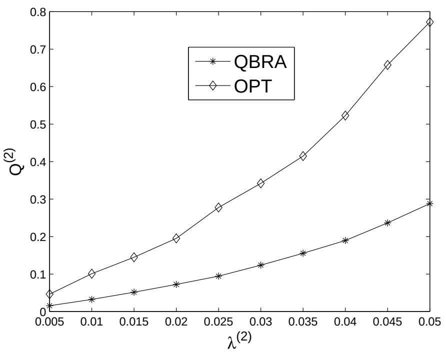

(b). Route 2: $Q^{(2)} = Q_1^{(2)} + Q_2^{(2)}$.

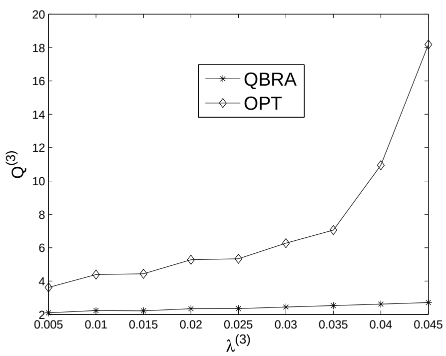

(c). Route 3: $Q^{(3)} = Q_1^{(3)} + Q_2^{(3)} + Q_3^{(3)}$.

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

we show simulation results for the system with exogenous arrivals, comparing queueing performance of QBRA with performance of optimization based algorithms (let's refer to them as OPT), such as that in [9].

First, we want to emphasize that, even when QBRA and OPT 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 OPT are oblivious of the real queues; QBRA can be implemented by nodes exchanging queueing information within local neighborhoods, while OPT 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: OPT 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 need not know or estimate input rates and ensures stability "automatically" (when feasible).

We simulate a system with exogenous (i.i.d. Poisson) arrivals with equal rates for all flows, $\lambda^{(1)} = \lambda^{(2)} = \lambda^{(3)}$ and scale them 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 of the queue lengths, except if there is 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 law is proportional to the end-to-end queueing delay): $Q^{(r)} = \sum_l Q_l^{(r)}$. Figure 2 compares $Q^{(1)}, Q^{(2)}, Q^{(3)}$ under the QBRA and OPT. It shows that the average queues under QBRA are significantly lower than the optimization-based algorithm. An intuitive explanation of this is that QBRA "better adapts" to the current queue length in the network.

References

[1] N. Abramson. The ALOHA system - another alternative for computer communications. In *Proceedings of AFIPS Conference*, volume 37, pages 281–285, 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*, March 2006.

[3] IEEE. Wireless lan medium access control (mac) and physical layer (phy) specificationsieee standard 802.11., June 2006. IEEE Standard 802.11.

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

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

[6] 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, 2000.

[7] A. Stolyar. Dynamic distributed scheduling in random access networks, 2005. Submitted.

[8] 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.

[9] 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*, 2005.