Optimal Throughput Allocation in General Random-Access Networks

Piyush Gupta
Bell Laboratories, Lucent Technologies
600 Mountain Ave., 2C-374
Murray Hill, NJ 07974
Email: pgupta@research.bell-labs.com

Alexander L. Stolyar
Bell Laboratories, Lucent Technologies
600 Mountain Ave., 2C-322
Murray Hill, NJ 07974
Email: stolyar@research.bell-labs.com

***Abstract*— We consider a model for random-access communication in networks of arbitrary topology. We characterize the efficient (Pareto) boundary of the network throughput region as the family of solutions optimizing weighted proportional fairness objective, parameterized by link *weights*. Based on this characterization we propose a general distributed scheme that uses *dynamic* link weights to "move" the link-throughput allocation within the Pareto boundary to a desired point optimizing a specific objective.**

As a specific application of the general scheme, we propose an algorithm seeking to optimize weighted proportional fairness objective *subject to minimum link-throughput constraints*. We study asymptotic behavior of the algorithm and show that link throughputs converge to optimal values as long as link dynamic weights converge. Finally, we present simulation experiments that show good performance of the algorithm.

I. Introduction

Next-generation wireless networks are likely to have a more decentralized architecture than the current cellular networks. For instance, in the emerging pico-cell architectures, the base station may not continue to perform the role of a central coordinating agent for the uplink access of user terminals. Such a decentralized architecture is already employed in 802.11-based wireless networks.

An important issue in such networks is that of scheduling. Due to the decentralized control constraints, a natural approach to consider is random-access communication, as in the slotted Aloha wireless LAN and 802.11 systems. It is well known, however, that the multi-user contention for channel access, if not well regulated, can lead to significant throughput degradation even in wireless LAN random-access systems ([3] and references therein), where any two concurrent transmissions interfere with each other. The problem gets substantially more aggravated in the networks with more general interference structure (such as those arising in multi-hop communication), where there are pronounced hidden-node and exposed-node issues. Hence, it becomes important to determine the optimal throughput that can be achieved in such more general random-access systems and to devise distributed control schemes that can operate the network close to it.

To address these problems, we consider a general model for random-access communication in networks of arbitrary topology. Briefly, the model is as follows (formal description in Section II). Consider a network with a finite set of nodes $\mathcal{N}$. Each node $n \in \mathcal{N}$ has one or several outgoing links (n, m) to a subset of other nodes $m \in \mathcal{D}_n \subset \mathcal{N}$. In each time slot, node n accesses the channel with probability p_n, and it chooses only one of its outgoing links to transmit on with probabilities p_{nm}/p_n. Each node transmits independently of other nodes. A transmission by node n interferes with and "erases" any simultaneous packet reception by any node k within a subset $\mathcal{N}_n \supset \mathcal{D}_n$. A transmission on a link (n, m) is successful if it is not erased by any other simultaneous transmission. The set of link throughputs $\mu = \{\mu_{nm}\}$ is thus a function of the set of access probabilities $p = \{p_{nm}\}$. This model is a generalization of the classical slotted Aloha system, proposed in [1] and analyzed in [8] (among others). It also generalizes recently studied models for general topology networks in [6], [12], where it is assumed that $k \in \mathcal{N}_n$ if and only if $n \in \mathcal{N}_k$. The model of this paper is also related to another general interference model of [4] - the main difference is that in [4] erasures due to interference occur with certain probabilities, not necessarily equal to 1 or 0.

As mentioned before, we seek to devise *distributed* control algorithms that achieve *efficient* throughput allocations μ. "Efficient" naturally means that allocation μ is (or close to) Pareto optimal, i.e., it is such that it cannot be improved upon.

Main Contributions

- For the general random-access network model above, we characterize the Pareto boundary M^* of the throughput region (under very mild additional condition), as a set of μ optimizing the following Weighted Proportional Fairness (WPF) objective:

$$\max_{\mu \in M} \sum_{(n,m)} w_{nm} \log \mu_{nm}, \qquad (1)$$

  for all possible sets of positive link "weights" $\{w_{nm}\}$. This characterization generalizes that given in [8], [9] for the classical Slotted Aloha, and is similar to that obtained in [4] (for a different model).
- Just as the characterization the Pareto boundary for the classical slotted Aloha model has provided many key insights utilized in the design of practical random-access systems, the above characterization of the Pareto boundary M^* leads us the following distributed procedure

for achieving a specific system objective. Link weights w_{nm} are used as *dynamic* control parameters; nodes *dynamically adjust* their link weights w_{nm} according to their own "satisfaction" with observed link throughputs; the link weights (in fact, only certain aggregates of them) are shared among neighboring nodes; nodes dynamically set their access probabilities to those optimal for the WPF objective with the current link weights. As a result, the set of link throughputs μ "moves" to a desired point while staying within the Pareto boundary. In other words, although WPF (with fixed weights) in itself is a very common and useful resource allocation objective in communication networks (cf. [7]), it can also be used as a "tool" for efficient throughput allocation in random-access networks for perhaps very different objectives.

- We apply the above general approach to the specific problem of achieving weighted proportional fairness *subject to minimum link throughput constraints*:

$$\max_{\mu\in M}\sum_{(n,m)}\alpha_{nm}\log\mu_{nm},\ \text{s.t. } \mu_{nm}\ge r_{nm},\forall(n,m), \quad (2)$$

  where $\alpha_{nm} > 0$ and $r_{nm} \ge 0$ are fixed parameters. We propose an algorithm utilizing very simple *token counter* mechanism, similar to that in [2], [10], that dynamically increases the dynamic weight w_{nm} of the links not achieving their required minimum throughput, so that eventually, in steady state, they do achieve it.
- We study the dynamics of token counters and link throughputs under the above algorithm, and prove its optimality in the sense that, if the token counters converge, then the link throughputs converge to the unique solution of problem (2). Finally, we provide simulation result which show good performance of the algorithm.

The rest of the paper is organized as follows. The formal model is described in Section II. The optimal solution to the WPF objective (with fixed weights) is in Section III. Section IV contains characterization of the throughput region, including smoothness properties of the Pareto boundary. In Section V we describe the distributed algorithm for problem (2). Section VI contains asymptotic analysis of the algorithm (as one of its parameters becomes small), and proves its optimality (in the sense described above). The simulation experiments are discussed in Section VII.

Basic Notation. We use the notations R, R_+ and R_{++} for the sets of real, real non-negative and real positive numbers, respectively. Corresponding I-times product spaces are denoted R^I, R^I_+, and R^I_{++}. The space R^I is viewed as a standard vector-space, with elements $x\in R^I$ being row-vectors $x=(x_1,\dots,x_I)$, and with Eucleadian metric induced by the norm $\|x\| \doteq \left[\sum_i x_i^2\right]^{1/2}$. Vector equalities and inequalities are understood componentwise.

II. The Model

Our model is as follows. (It is a generalization of the model of [6], [12].) The system consists of a finite set $\mathcal{N}=\{1,2,\dots,N\}$ of *nodes*, and operates in discrete time, with time slots indexed by $t=0,1,2,\dots$. Let $\mathcal{D}_n\subseteq\mathcal{N}\setminus n$ denote the subset of nodes to which node n has data to send. A node n at any time t may attempt transmission of one unit of data (say, data packet) to one of the nodes $m\in\mathcal{D}_n$. When this happens, we say that node n makes transmission attempt on the *link* (n,m). We will denote by

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

the set of all system links, and by I its cardinality (i.e., the total number of links).

We assume that a node cannot simultaneously (i.e., within the same slot) transmit on two or more different links. The interference between simultaneous transmissions in the network has the following structure. If a node transmits in a slot, any simultaneous attempt to transmit to this node will fail. If there are two or more simultaneous transmissions to a node, they all collide and fail. Any transmission attempt by node n will interfere with and "erase" any attempt to receive a message at any of the nodes within some subset of $\mathcal{N}$, denoted by $\mathcal{N}_n$. (The model of [6], [12] additionally assumes that $m\in\mathcal{N}_n$ implies $n\in\mathcal{N}_m$.) Given the above assumptions, $\mathcal{D}_n\subseteq\mathcal{N}_n$. (In other words, a transmission attempt by node n may interfere with receiving at more nodes than it actually sends traffic to.) Also, because a node n transmission makes simultaneous successful receiving impossible, $n\in\mathcal{N}_n$, for all n.

Consider the following "Slotted Aloha-type" random access strategy. Each node n in each time slot transmits with probability p_n, independently of other nodes and of the past history. And when node n does transmit, it chooses a particular link to transmit on, among the links $(n,m),\ m\in\mathcal{D}_n$, also randomly, with probabilities p_{nm}/p_n summing up to 1, that is

$$\sum_{m\in\mathcal{D}_n}p_{nm}=p_n. \quad (3)$$

Given this strategy, the average throughputs on the network links are given by

$$\mu_{nm}=p_{nm}\prod_{k:\ m\in\mathcal{N}_k,\ k\neq n}(1-p_k). \quad (4)$$

The dependence of the set (vector) of throughputs $\mu=(\mu_{nm},\ (n,m)\in\mathcal{I})\in R^I_+$ on the set (vector) of access probabilities

$$p\in\mathcal{P}\doteq\{(p_{nm},\ (n,m)\in\mathcal{I})\in[0,1]^I\mid (3)\text{ holds}\},$$

given by (4), will be denoted by $\mu(p)$. Clearly, function $\mu(p)$ is continuous.

III. Optimal Solution for the Weighted Proportional Fairness Objective

The following Theorem 1 is a generalization of the corresponding result in [6], in that it applies to a more general model and optimization objective. (It also generalizes some of the results of [4].) The theorem shows that the problem of choosing access probabilities optimizing the *weighted proportional fairness* objective is relatively easy to solve, and it serves as a starting point for the development in this paper.

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

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

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

Theorem 1: For arbitrary set of positive weights $w = \{w_{nm},\ (n,m) \in \mathcal{I}\} \in R^I_{++}$, there exists a unique set of access probabilities $p \in \mathcal{P}$ that maximizes the function

$$F = \sum_{(n,m)\in\mathcal{I}} w_{nm} \log \mu_{nm}. \quad (5)$$

The optimal p is given by:

$$p_{nm} = \frac{w_{nm}}{\sum_{(\ell,k)\in\mathcal{S}_n} w_{\ell k}}. \quad (6)$$

Remark 1. Expression (6) can be equivalently rewritten as

$$p_{nm} = \frac{w_{nm}}{\sum_{m\in\mathcal{N}_n} W^{in}_m}, \quad (7)$$

where

$$W^{in}_m \doteq \sum_{\ell:\ m\in\mathcal{D}_\ell} w_{\ell m} \quad (8)$$

is the sum of the weights of all links "incoming" to node m; we will call W^{in}_m the *incoming weight* of node m.

Proof of Theorem 1. Consider a fixed node n. Suppose first that the set

$$\mathcal{S}^-_n \doteq \{(\ell,k) \in \mathcal{S}_n \mid \ell \neq n\}$$

is non-empty. (In other words, node n's transmissions interfere with transmissions on at least one link not originating at n.) In this case, any p maximizing F must be such that $0 < p_{nm} \leq p_n < 1$ for all $m \in \mathcal{D}_n$. Then, if we substitute (4) and (3) into (5), we see that

$$\frac{\partial F}{\partial p_{nm}} = \frac{w_{nm}}{p_{nm}} - \sum_{(\ell,k)\in\mathcal{S}^-_n} \frac{w_{\ell k}}{1-p_n} = 0,$$

which yields

$$p_{nm} = (1-p_n)\frac{w_{nm}}{\sum_{(\ell,k)\in\mathcal{S}^-_n} w_{\ell k}}. \quad (9)$$

Summing up (9) over $m \in \mathcal{D}_n$, we obtain an equation for p_n, whose solution is

$$p_n = \frac{\sum_{m\in\mathcal{D}_n} w_{nm}}{\sum_{m\in\mathcal{D}_n} w_{nm} + \sum_{(\ell,k)\in\mathcal{S}^-_n} w_{\ell k}} = \frac{\sum_{m\in\mathcal{D}_n} w_{nm}}{\sum_{(\ell,k)\in\mathcal{S}_n} w_{\ell k}}. \quad (10)$$

Expressions (10) and (9) give (6). In the case when $\mathcal{S}^-_n = \emptyset$, it is easy to see that the access probabilities $p_{nm},\ m \in \mathcal{D}_n$, must maximize $\sum_m w_{nm} \log p_{nm}$ subject to $\sum_m p_{nm} \leq 1$. The unique solution is

$$p_{nm} = \frac{w_{nm}}{\sum_{k\in\mathcal{D}_n} w_{nk}}.$$

However, $\mathcal{S}^-_n = \emptyset$ means that $\{(n,k) \mid k \in \mathcal{D}_n\} = \mathcal{S}_n$, and thus expression (6) is still valid. ∎

The dependence of the set (vector) of access probabilities $p \in \mathcal{P}$ on the set (vector) of positive link weights $w \in R^I_{++}$, given by (6), will be denoted by $p(w)$. (Clearly, $p(w)$ is invariant with respect to scaling of w by a positive constant.)

IV. System Throughput Region Characterization

From this point on in the paper, for brevity, we sometimes denote links $(n,m) \in \mathcal{I}$ by a single index i, j, etc.

We define the system *throughput region* M as the set of all non-negative vectors, which can be majorized by vectors of the form $\mu(p)$, namely,

$$M \doteq \{\mu' \in [0,1]^I \mid \mu' \leq \mu(p) \text{ for some } p \in \mathcal{P}\}. \quad (11)$$

We denote by

$$M^* \doteq \{\mu^* \in M \mid \mu^* \leq \mu' \in M \text{ implies } \mu' = \mu^*\} \quad (12)$$

the subset of maximal elements of M, which can be called the *Pareto* boundary of M. Characterizing boundary M^* is the main focus of this section. We denote by

$$M^*_{++} \doteq M^* \cap R^I_{++} \quad (13)$$

the subset of Pareto boundary M^* consisting of vectors with all strictly positive components.

The following proposition describes basic properties of the throughput region M^*. We omit the straightforward proof.

Proposition 1: (i) Throughput region M is a compact set.

(ii) Set M^* is non-empty. For any $\mu^* \in M^*$ there exists $p \in \mathcal{P}$ such that $\mu^* = \mu(p)$.

It follows from Theorem 1 that for any $w^* \in R^I_{++}$, $\mu^* = \mu(p(w^*)) \in M^*_{++}$. The natural question is whether or not the converse is true, namely, that for any $\mu^* \in M^*_{++}$ we can find w^* such that $\mu^* = \mu(p(w^*))$. The answer is basically *yes*, under a mild additional condition, as we show below in Theorem 2.

Let U denote the system log-throughput region. More precisely,

$$U \doteq \{\log x \mid x \in M \cap R^I_{++}\}, \quad (14)$$

where here and below $\log$ applied to a vector is understood component-wise. The Pareto boundary of U is

$$U^* \doteq \{\log x \mid x \in M^* \cap R^I_{++} \equiv M^*_{++}\}. \quad (15)$$

Lemma 1: The system log-throughput region U is a closed convex subset of the negative orthant of R^I.

Proof. It is easy to observe that, for any link $i \in \mathcal{I}$, $\log \mu_i(p)$ is a concave scalar function of the set of access probabilities p. Then, for any $\mu^{(1)} = \mu(p^{(1)}) \in M \cap R^I_{++}$ and $\mu^{(2)} = \mu(p^{(2)}) \in M \cap R^I_{++}$, a convex combination of $\log \mu^{(1)}$ and $\log \mu^{(2)}$ is

$$u = \alpha_1 \log \mu^{(1)} + \alpha_2 \log \mu^{(2)} \leq \log \mu(\alpha_1 p^{(1)} + \alpha_2 p^{(2)}) \in U.$$

This means that $u \in U$ as well. Since any $u \in U$ is dominated by $u' = \log \mu(p') \in U$ for some p', the convexity of U follows. Region U is closed because M is closed, and $\log$ is a continuous mapping. ∎

Theorem 2, presented just below, characterizes the Pareto boundary of the throughput region. This result is analogous to the results of Sections 3.4-3.5 of [4], which apply to a closely related - but different - model. In particular, Theorem 2 generalizes Theorem 6 of [4].

Consider the directed graph with vertices being links $i \in \mathcal{I}$, and the edge from $i = (n,m)$ to j existing if and only if $j \in \mathcal{S}_n \setminus i$. We will call this graph a *link dependence graph*. (In the case when there is at most one link originating from each node, this graph could be called "interference graph" - the term used in [4].)

Since function $p(w)$ is invariant with respect to scaling of w, we can restrict the domain R^I_{++} of $p(w)$ to the normalized set $B \doteq \{w \in R^I_{++} \mid \sum_i w_i = 1\}$.

Theorem 2: Suppose, the link dependence graph is strongly connected. (There is a directed path from any vertex to any other.) Then, function $\mu(p(w))$ defines a homeomorphism (mutually continuous one-to-one mapping) between B and M^*_{++}. Moreover, M^*_{++} is a smooth $(I-1)$-dimensional surface.

Proof of Theorem 2. The outline of the proof that $\mu(p(w))$ is a homeomorphism is as follows. For any $w \in B$, $\mu(p(w)) \in M^*_{++}$. Then, we establish the following sequence of assertions, for a fixed $\mu^* \in M^*_{++}$.

Assertion 1. There exists $w^* \in R^I_{++}$ such that $\mu^* = \mu(p(w^*))$.

Assertion 2. Vector $p(w^*)$ is the unique vector $p \in \mathcal{P}$ solving equation $\mu^* = \mu(p)$.

Assertion 3. Vector w^* is the unique vector $w \in B$ solving equation $\mu^* = \mu(p(w^*))$.

Assertion 4. If $\mu' \to \mu^*$, then $w' \to w^*$, where $\mu' \in M^*_{++}$ and $\mu' = \mu(p(w'))$.

Due to space limitation, we do not give the detailed proof of these assertions, and of the smoothness of M^*_{++}. (All proofs are given in [5].) ∎

V. Providing Minimum Link Throughput Guarantees: A Distributed Algorithm

Suppose we want an efficient "distributed" random access algorithm that satisfies certain minimum link throughput requirements whenever this is feasible at all. It is also desirable that the "leftover" system capacity, after satisfying the minimum throughput constraints, is allocated in a "fair" fashion. We will combine and specify this two objectives as follows. Suppose, a weight $\alpha_i > 0$ and a minimum throughput requirement $r_i \geq 0$ is given for each link $i \in \mathcal{I}$, that is, there are two parameter vectors $\alpha = \{\alpha_i,\ i \in \mathcal{I}\} \in R^I_{++}$ and $r = \{r_i,\ i \in \mathcal{I}\} \in R^I_{+}$. We want a distributed random access algorithm such that, in steady state, the link throughput allocation μ^* solves the following optimization problem:

$$\max_{x \in M} \sum_i \alpha_i \log x_i \tag{16}$$

subject to

$$x \geq r. \tag{17}$$

Equivalently, in terms of log-throughput region U, we seek μ^* such that $u^* = \log \mu^*$ solves the problem:

$$\max_{u \in U} \sum_i \alpha_i u_i \tag{18}$$

subject to

$$u \geq \log r. \tag{19}$$

Note that if μ^* is a solution to (16)-(17) and it has the form

$$\mu^* \in \arg\max_{x \in M} \sum_i \alpha^*_i \log x_i \tag{20}$$

for some $\{\alpha^*_i,\ i \in \mathcal{I}\} \in R^I_{++}$, then $\mu^* \in R^I_{++}$ and this solution is unique. This and Theorem 2 easily imply that when the link dependence graph is strongly connected, the solution μ^* of (16)-(17) (if any) is unique.

The algorithm we propose is as follows. As before, let $t = 0,1,2,\ldots$ denote a time slot.

(A) Each node n, maintains a "token counter" (token queue length) $Q_{nm}(t)$ for each of its outgoing links (n,m), which is updated according to the following rule:

$$Q_{nm}(t+1) = Q_{nm}(t) + r_{nm} - h_{nm}(t)\ ,\ t = 0,1,2,\ldots\ , \tag{21}$$

where $h_{nm}(t) = 1$ if there was a successful transmission on link (n,m) in slot t, and $h_{nm}(t) = 0$ otherwise. (We will use vector notation $Q(t) = (Q_i(t),\ i \in \mathcal{I})$.)

(B) Each node n, for each of its outgoing links (n,m), calculates dynamic weight $w_{nm}(t) = \alpha_{nm} + \beta Q_{nm}(t)$, where $\beta > 0$ is some (typically small) parameter, and sets its access probabilities in slot t according the expression (7).

Implementation considerations. Of course, the big question is: How can a node n "know" the incoming weights W^{in}_m (defined in (8)) of the nodes $m \in \mathcal{N}_n$, i.e. nodes it interferes with? The form of (7) allows (at least in principle) the following natural procedure. Each transmission on any link (ℓ,k) contains the piggy-backed current dynamic weight $w_{\ell k}(t)$ of the link. Thus, each node m can maintain an estimate of its incoming weight W^{in}_m. Each node m periodically "broadcasts" its W^{in}_m, so it can be "heard" by all the nodes whose transmissions may interfere with reception by node m. (In particular, W^{in}_m can be piggy-backed into transmissions from node m.) Each node n listens to the broadcast messages described above, which allows it to estimate the sum in the denominator of (7).

VI. Asymptotic behavior of the algorithm with small parameter β

In this section we study the dynamics of user throughputs and token counters, under the algorithm described in Section V, when parameter $\beta > 0$ is small. Namely, we consider an asymptotic regime such that β converges to 0. We study the dynamics of *fluid sample paths* (FSP), which are (roughly speaking) possible trajectories $q(t)$ of a random process that is a limit of the process $\beta Q(t/\beta)$ as $\beta \to 0$. (In other words, trajectories $q(t)$ "approximate" the process $Q(t)$ scaled down by factor β, and with $1/\beta$ time speed-up.) The main result of this section (Theorem 3) basically says that if FSP is such that $q(t)$ converges to some finite vector q^* as $t \to \infty$, then link throughputs converge to the unique solution to the problem (16)-(17).

Remark 2. A stronger result would be to prove that the convergence of $q(t)$ in fact holds as long as problem (16)-(17) is feasible; this would prove the asymptotic optimality of the algorithm, as it is done, for example, for the Greedy Primal-Dual (GPD) [10] algorithm, for a *different model*. Proving asymptotic optimality of the algorithm of this paper may be a subject of future work. We note that, since problem (18)-(19) is convex, the GPD algorithm can in principle be applied (and be provably optimal), if we would "work" with logarithms of the throughputs; this however would require that we measure throughputs over longer time intervals (and thus updates token counters less frequently), which would result in a much "slower" algorithm.

Remark 3. It is shown in [11] that the algorithm of this paper, *with any fixed parameter* β, ensures that the users will receive the desired minimum throughputs as long as the constraint (17) is feasible. However, the results of [11] do not address the utility maximization objective (16).

We now define the asymptotic regime and an FSP. Let us denote by

$$F_i(t) = \sum_{s=0}^{t-1} h_i(s),\ t = 0, 1, 2, \ldots,\quad i \in \mathcal{I},$$

the total number of successful transmissions on link i by (and excluding) time t, and denote $F(t) = (F_i(t),\ i \in \mathcal{I})$. We extend the time domain of functions $F(t)$ and $Q(t)$ to all real $t \geq 0$ by adopting the convention that they are constant within each time slot $[t, t+1)$ for all integer $t \geq 0$.

Consider a sequence $\{\beta\}$ of positive values of β, converging to 0. For each β, let $(F^\beta(\cdot), Q^\beta(\cdot))$ be a realization of the corresponding random process, with some fixed initial $Q^\beta(0)$. Assume that the sequence of realizations is such that a functional law of large numbers (FLLN) condition holds for the process governing transmission attempt decisions by node n at different times t, given its "current" (depending on t) set of access probabilities. (The precise condition is given in [5].)

Consider the following rescaled trajectory for each β:

$$(f^\beta = (f^\beta(t),\ t \geq 0),\ q^\beta = (q^\beta(t),\ t \geq 0)),$$

where $f^\beta(t) = \beta F^\beta(t/\beta)$ and $q^\beta(t) = \beta Q^\beta(t/\beta)$.

Definition: A pair of vector-functions $(f = (f(t),\ t \geq 0),\ q = (q(t),\ t \geq 0))$ is called a *fluid sample path* (FSP), if the uniform on compact sets (u.o.c.) convergence $(f^\beta, q^\beta) \to (f, q)$ holds for at least one sequence $\{\beta\}$ and the corresponding sequence (f^β, q^β) of scaled trajectories, as defined above.

Theorem 3: Suppose an FSP (f, q) is such that

$$q(t) \to q^* \in R_+^I \quad \text{as } t \to \infty.$$

Then, the problem (16)-(17) is feasible, with $\mu^* = \mu(p(\alpha + q^*))$ being its unique optimal solution, and $(d/dt)f(t) \to \mu^*$.

To prove Theorem 3, we will first describe the basic FSP properties in Lemma 2.

Lemma 2: The family of fluid sample paths has the following properties.

(i) All component functions $f_i(t),\ t \geq 0$, and $q_i(t),\ t \geq 0$, are Lipschitz continuous, uniformly across all FSPs. Consequently, any FSP is such that proper derivatives $f_i'(t)$ and $q_i'(t)$ exist for almost all $t \geq 0$ (with respect to Lebesgue measure).
(ii) "Shift property." If (f, q) is an FSP, then for any $d \geq 0$, $(\theta_d f, \theta_d q)$ is also an FSP, where

$$[\theta_d f](t) = f(t+d) - f(d),\ [\theta_d q](t) = q(t+d), \quad t \geq 0.$$

(iii) "Compactness." If a sequence of FSPs $(f^{(j)}, q^{(j)}) \to (f, q)$ uniformly on compact sets as $j \to \infty$, then (f, q) is also an FSP.
(iv) For any FSP, for almost all $t \geq 0$ we have:

$$f'(t) = \mu(p(\alpha + q(t))), \tag{22}$$

and, for each $i \in \mathcal{I}$,

$$q_i'(t) = \begin{cases} r_i - f_i'(t) & \text{if } q_i(t) > 0, \\ \max\{r_i - f_i'(t), 0\} & \text{if } q_i(t) = 0. \end{cases} \tag{23}$$

(Consequently, since $f(t)$ is Lipschitz and $\mu(p(w))$ is continuous, the derivative $f'(t)$ exists for all $t > 0$.)

Proof. See [5]. ∎

Proof of Theorem 3. Using shift and compactness properties, from the FSP as in the theorem statement, it is easy to see that the "stationary" trajectory, given by

$$q(t) \equiv q^*, \quad f'(t) \equiv \mu^* = \mu(p(\alpha + q^*)) \tag{24}$$

is also an FSP. We immediately see that vector r cannot lie outside region M, for otherwise $r_i > \mu_i^*$ for at least one i, implying (by (23)) that $q_i(t) \to \infty$, which contradicts (24).

If we use notation $u^* = \log \mu^*$, then, by the definition of function $\mu(p(\cdot))$ (see (5)),

$$u^* \in \arg\max_{u \in U} \sum_i (\alpha_i + q_i^*) u_i. \tag{25}$$

If we specialize (23) for the stationary FSP defined above, we see that $q_i^* > 0$ implies $\mu_i^* = r_i$, or, equivalently,

$$q_i^*(u_i^* - \log r_i) = 0 \text{ for all } i. \tag{26}$$

If we view q_i^* as Lagrange multipliers for the constrained optimization problem (18)-(19), we see that, by Kuhn-Tucker theorem, (25) and (26) imply that u^* is a solution to that problem. The uniqueness of solution u^* follows from the representation (25), as noted in the remark containing (20). ∎

VII. Simulation results

We now present simulation results for the algorithm introduced in Section V. Throughout this section we always assume that weights $\alpha_i = 1$ (see (16)) for all links, and so the system objective is to maximize the sum of the logarithms of link throughputs, subject to minimum link throughput constraints (17). The parameter $\beta = 0.001$ is same throughout all experiments.

We start with a simple 3-node network, shown in Figure 1. A bi-directional link on the figure, say between nodes 1 and 2, means that both $(1, 2)$ and $(2, 1)$ are communication (data transmission) links of the system. There is no interference between nodes 2 and 3. (That is, $3 \notin \mathcal{N}_2$ and $2 \notin \mathcal{N}_3$.) For this network, expressions (8) and (7) specialize to:

$$W_1^{in} = w_{21} + w_{31}, \; W_2^{in} = w_{12}, \; W_3^{in} = w_{13},$$

$$p_{1m} = \frac{w_{1m}}{W_1^{in} + W_2^{in} + W_3^{in}}, \; p_{m1} = \frac{w_{m1}}{W_1^{in} + W_m^{in}}, \; m = 2, 3.$$

Table I shows steady state link throughputs (after they converge) for two cases. The first is the "baseline" case when we do not impose any minimum throughput requirements, i.e., $r_{nm} = 0$ for all links. In the second case, we introduce minimum throughput requirement $r_{2,1} = 1/7$ for link $(2, 1)$, and leave $r_{nm} = 0$ for all other links. We see that, in the second case, the algorithm indeed "lifts" the throughput of link $(2, 1)$ to the desired minimum level (at the expense of the throughputs on the other links, of course.) The token counters (and therefore the dynamic weights) in the second scenario indeed "converge" (see [5]), up to some inevitable "jitter" since β is finite; this guarantees (by Theorem 3) that the throughput allocation is indeed optimal.

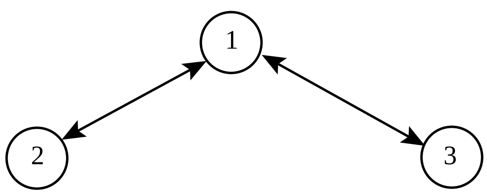

Fig. 1. 3-node system.

	all $r_{n,m} = 0$	$r_{2,1} = 1/7 (= 0.1429)$
$\mu_{1,2}$	0.1649	0.1437
$\mu_{2,1}$	0.1115	0.1432
$\mu_{1,3}$	0.1677	0.1609
$\mu_{3,1}$	0.1114	0.1023

TABLE I

3-NODE SYSTEM. STEADY STATE LINK THROUGHPUTS.

Next, we consider a more complex 10-node network, shown in Figure 2. As before, a solid bi-directional link means presence of a communication link in both directions. A bi-directional dashed link means mutual interference between nodes. (For example, $4 \in \mathcal{N}_9$ and $9 \in \mathcal{N}_4$. Recall that also, if a communication link (n, m) exists, then, by our definitions, n causes interference to m; for example, $10 \in \mathcal{N}_4$.) As for the 3-node network, we consider two cases: "baseline," with $r_{nm} = 0$ for all links, and the second case where we have minimum throughput requirement for one of the links, namely $r_{5,9} = 0.1$ for link $(5, 9)$. Steady state link throughputs for both cases are shown in Table II. We see that the throughput of link $(5, 9)$ is indeed "lifted" to approximately the required level. (An interesting - although not very surprising - observation is that, in systems of general topology, introducing minimum throughput requirement on one of the links does *not* necessarily result in the decrease of the throughputs on *all* other links.) Again, simulation shows (see [5]) that the token counters indeed "converge," which guarantees optimality of the the throughput allocation.

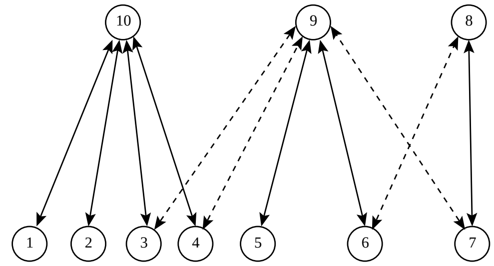

Fig. 2. 10-node system.

	all $r_{n,m} = 0$	$r_{5,9} = 0.1$
$\mu_{1,10}$	0.0569	0.0492
$\mu_{10,1}$	0.0901	0.0896
$\mu_{2,10}$	0.0564	0.0490
$\mu_{10,2}$	0.0894	0.0891
$\mu_{3,10}$	0.0359	0.0287
$\mu_{10,3}$	0.0689	0.0744
$\mu_{4,10}$	0.0380	0.0290
$\mu_{10,4}$	0.0670	0.0742
$\mu_{5,9}$	0.0412	0.0986
$\mu_{9,5}$	0.0691	0.0492
$\mu_{6,9}$	0.0876	0.0687
$\mu_{9,6}$	0.0716	0.0668
$\mu_{7,8}$	0.1245	0.1048
$\mu_{8,7}$	0.1817	0.2042

TABLE II

10-NODE SYSTEM. STEADY STATE LINK THROUGHPUTS.

Acknowledgement. We are grateful to Yuliy Baryshnikov for pointing out that the convexity of log-throughput region can be directly observed, as in the proof of Lemma 1.

REFERENCES

[1] N. Abramson. The ALOHA system – Another alternative for computer communications. *Proc. AFIPS Conf.*, Vol. 37, (1970), pp. 281-285.

[2] M. Andrews, L. Qian, A. L. Stolyar. Optimal Utility Based Multi-User Throughput Allocation subject to Throughput Constraints. *Proceeding of INFOCOM'2005*, Miami, March 13-17, 2005.

[3] P. Gupta, Y. Sankarasubramaniam, A. L. Stolyar. Random-Access Scheduling with Service Differentiation in Wireless Networks. *Proceeding of INFOCOM'2005*, Miami, March 13-17, 2005.

[4] P. Gupta, A. L. Stolyar. Throughput Region of Random Access Networks of General Topology. Bell Labs Technical Memo, 2005. Submitted.

[5] P. Gupta, A. L. Stolyar. Optimal Throughput Allocation in Random Access Networks of General Topology. Bell Labs Technical Memo, 2005.

[6] K. Kar, S. Sarkar, L. Tassiulas. Achieving Proportionally Fair Rates using Local Information in Aloha Networks. *IEEE Trans. Autom. Control*, Vol. 49, (2004), No. 10, pp. 1858–1862.

[7] F. P. Kelly. Mathematical Modelling of the Internet. In: Engquist B., Schmidt W. (eds). *Mathematics Unlimited - 2001 and Beyond.* Springer, Berlin, 2001, pp. 685-702.

[8] J. Massey and P. Mathys. The collision channel without feedback. *IEEE Trans. Inform. Theory,* Vol. IT-31, (1985), no. 2, pp. 192–204.

[9] K. A. Post. Convexity of the Nonachievable Rate Region for the collision channel without feedback. *IEEE Trans. Inform. Theory,* Vol. IT-31, (1985), No. 2, pp. 205–206.

[10] A. L. Stolyar. Maximizing Queueing Network Utility subject to Stability: Greedy Primal-Dual Algorithm. *Queueing Systems*, 2005, Vol.50, No.4, pp.401-457.

[11] A. L. Stolyar. Dynamic Distributed Scheduling in Random Access Networks. Bell Labs Technical Memo, 2005. Submitted.

[12] X. Wang and K. Kar. Distributed Algorithms for Max-min Fair Rate Allocation in Aloha Networks. *Proceedings of the 42nd Annual Allerton Conference*, Urbana-Champaign, 2004.