Throughput Region of Random Access Networks of General Topology

Piyush Gupta and Alexander L. Stolyar

Abstract

A random access model is introduced and studied, which is a generalization of the classical Slotted Aloha model. Unlike in the Slotted Aloha, where two or more simultaneous transmissions on any subset of links collide and "erase" each other, a quite general interference structure is considered, where transmission on link i erases a simultaneous transmission on link j with some fixed probability ϕ_{ij}. In particular, it is allowed that $\phi_{ij} \neq \phi_{ji}$, which captures possible asymmetric interference in real - most notably wireless - communication networks.

Results characterizing the maximum achievable link throughput region and its Pareto boundary, are derived. In some cases, the Pareto boundary characterization is almost as simple and explicit as that derived in prior work for the classical Slotted Aloha system.

Index Terms

Collision channel, communication networks, interfence, random access, slotted Aloha, throughput region, weighted proportional fairness, wireless.

I. Introduction

In this paper we study models of random access communication networks of general topology. In particular, we introduce and analyze a model, which is a generalization of the classical Slotted Aloha random access system (cf. [11] for the Slotted Aloha, and [2] for the original Aloha system).

In the Slotted Aloha model, there are I "communication links" indexed $i = 1, \dots, I$. In each time slot, a transmission on link i is attempted with probability p_i. In case of simultaneous transmission attempts on two or more different links, all such transmissions "collide," and as a result, they all get "corrupted" and fail. (That is, all links "interfere" with each other.)

The model we focus on in this work is more general than Slotted Aloha in that it allows a much more general interference structure. Namely, a transmission on link i "corrupts", or "erases", a simultaneous transmission on link j with a certain probability ϕ_{ij}. In addition to capturing natural "spatial" features of real communication (especially wireless) networks, the model allows asymmetric interference in the sense that the cases when $\phi_{ij} \neq \phi_{ji}$ are included.

A fundamental question for our model (as well as for any random access model in general) is that of characterizing the system *throughput region* M, namely, the set of all achievable long-term average link throughput vectors $\mu = (\mu_1, \dots, \mu_N)$, which can be obtained by varying access probabilities p_i. In turn, the key further question is the characterization of the *Pareto* boundary M^*, consisting of vectors $\mu \in M$, which cannot be "improved upon." The throughput region of the Slotted Aloha system was fully characterized in [11], where it was shown that a throughput vector μ lies on the Pareto boundary M^* if and only if it corresponds to a set of access probabilities p_i summing up to 1. Besides its elegance, this result has proved to be very useful in practice. Indeed, it provides a "guiding principle" for the design of practical random access contention based schemes, in which "access rates"

P. Gupta and A.L.Stolyar are with Bell Labs, Alcatel-Lucent, Murray Hill, NJ, 07974 USA e-mails: {pgupta,stolyar}@research.bell-labs.com

Manuscript received Month NN, 2005; revised Month NN, 2011.

(roughly corresponding to access probabilities) are dynamically adapted. (See our paper [4] for an example of such a dynamic scheme, called RCMAC, and for an extensive review of related literature.)

In this paper we study properties of the throughput region M and its Pareto boundary M^* for the general model described above. The mathematical techniques used in [11] do not apply to our model and, consequently, we use a different approach. Our results provide different degree of characterization of M^*, depending on additional assumptions. In some special cases (which are still far more general than the Slotted Aloha) the characterization of M^* is almost as simple and explicit as that given in [11] for the Slotted Aloha. For example, suppose that each ϕ_{ij} is either 0 or 1 and the "interference graph" (with vertices $i = 1, \ldots, I$, and edges (i, j) existing if $\phi_{ij} = 1$) is strongly connected. Then we show that vector μ (with positive components) lies on the Pareto boundary M^* if and only if the access probabilities are given by

$$p_i(w) = \frac{w_i}{w_i + \sum_{j \in \mathcal{I}_i} w_j}, \tag{1}$$

where $w_i > 0$ are some weights assigned to the links, and $\mathcal{I}_i$ denotes the set of links to which link i causes interference. Thus, in this case, *any efficient throughput allocation* can be achieved if we assign each link a weight w_i and make sure that each link "knows" the weights of those of its "neighboring" links, to which it causes interference. By assigning different weights w_i one can "move" within the boundary M^* to try to optimize a specific system-wide objective.

A model which is a different ("node-centric") generalization of Slotted Aloha was studied in [9], [14]. (In this model the interference is of "0-1" type, which in terms of our model means that all ϕ_{ij} are 0 or 1.) In [9], the problem of achieving *proportional fairness*, that is $\max \sum_i \log \mu_i$, has been solved. (As demonstrated in [9], the technical advantage of the "sum-log" objective function is that the "multiplicative form" of dependence of each μ_i on the access probabilities p_j makes the optimization problem relatively easily tractable. In fact, our expressions of the type (1) arise as optimizers of the *weighted proportional fairness* problem, $\max \sum_i w_i \log \mu_i$, which serves as a starting point of our analysis.) Paper [14] provides solutions to the *max-min fairness* problem, $\max \min_i \mu_i$.

A model with probabilistic (rather that "0-1") interference structure was considered in [3]. Specifically, an infinite-user slotted aloha with multi-packet reception capability was analyzed, where the number of correctly received packets in a slot is a probabilistic function of the number of simultaneous transmitted packets. It was shown that this system is stable if the packet arrival rate is less than the (asymptotic) expected number of packets successfully received in a collision. This model, however, is very different from the one considered in this paper: in the former, the number of successfully received packets depends only on the number of simultaneous transmissions, irrespective of the transmitter-receiver-interferers topology, which is explicitly captured in the model considered here.

Finally, we note that our model, as well as all other models mentioned above, assumes "saturated queues," that is, roughly speaking, each link always has "data packets" to transmit. For the models with exogenously (randomly) arriving packets, which are queued until transmitted, see [1] and references therein.[1]

The rest of the paper is organized as follows. In Section II we introduce notations used throughout the paper. The model, the random-access strategy and its throughput region are defined in Sections III and IV. In Section V we discuss non-convexity of the throughput region. Section VI gives solution for the weighted proportional fair objective and describes its basic properties - this serves as a starting point for further analysis. The main results, characterizing Pareto boundary M^* and its smoothness properties, in the case of a strongly connected interference graph, are in Sections VII and VIII, respectively.

II. Notation and Conventions

We use the standard notations R and R_+ for the sets of real and real non-negative numbers, respectively; and the not quite standard R_{++} for the set of *strictly* positive real numbers. Corresponding I-times product spaces

[1]While this paper was under review (which took a very long time) there has been a large amount of new work on random access schemes. In particular, the "node-centric" model, closely related to the model of this paper, was further studied in [5], [13], [6]. A very active new line of research was originated by paper [7], where a class of Carrier-Sense/Collision-Avoidance schemes was proposed that achieves throughput optimality in a variety of contexts. We refer reader to [8] for a good review of recent advances in the analysis of random access algorithms.

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 $x = (x_1, \ldots, x_I)$. We write just 0 for the zero vector or its transpose. Vector inequalities are understood componentwise:

$$x \le (<) y \quad \Leftrightarrow x_i \le (<) y_i, \ \forall i.$$

The scalar product (dot-product) of $x, y \in R^I$, is

$$x \cdot y \doteq \sum_{i=1}^{I} x_i y_i \ ;$$

and the norm of x is

$$\|x\| \doteq \sqrt{x \cdot x} \ .$$

Sometimes, we use the following slightly abusive notations:

$$(x/y) \doteq (x_1/y_1, \ldots, x_I/y_I), \quad (1/y) \doteq (1/y_1, \ldots, 1/y_I),$$
$$\sqrt{x} \doteq (\sqrt{x_1}, \ldots, \sqrt{x_I}), \quad \log x \doteq (\log x_1, \ldots, \log x_I).$$

Matrix $diag(x)$ is the diagonal matrix, with diagonal elements given by the components of vector x. For a matrix J, J^T and $rank(J)$ denote its transpose and rank, respectively. For a vector-function $x = x(y)$ mapping R^I into itself, $(\partial x/\partial y)$ will denote its $I \times I$ Jacobi matrix $(\partial x_i/\partial y_j, \ i \in I, \ j \in I)$.

We denote by $\overline{A}$ the closure of a set A.

III. A Model of Random-Access Networks of General Topology

We consider a quite general model of a communication network, where simultaneous data transmissions on different links may interfere with each other. (Wireless networks serve as primary motivation and application of the model.) The interfence between links is modeled by assuming that a transmission on link i "erases" a *simultaneous* transmission (if any) on link j with a certain probability ϕ_{ij}, $0 \le \phi_{ij} \le 1$. Probability ϕ_{ij} is *not* necessarily equal to ϕ_{ji}. Thus, our model includes scenarios where link i interferes with link j, but not vice versa; such scenarios are *very common* is wireless mesh networks. We seek to optimize the vector $\mu = \{\mu_i\}$ of (useful) throughputs on the links. The optimization objectives may vary. Our primary goal will be to characterize the class of controls producing maximal (Pareto optimal) vectors μ. Then, optimizing throughputs to satisfy a specific objective is typically reduced to optimization within the class of Pareto optimal controls. We now introduce the model and notations more formally.

The system consists of a finite set $\mathcal{I} = \{1, 2, \ldots, I\}$ of *vertices*. (A vertex $i \in \mathcal{I}$ in this model typically corresponds to a data transmission link from a physical node n_1 to another physical node n_2 of a communication - for instance, wireless - network. That's why this model can be called "link-centric.") The system operates in discrete time, with time slots indexed by $t = 1, 2, \ldots$. Any vertex i at any time t may attempt transmission of one unit of data (say, data packet). When this happens, we say that vertex i is *active*, or makes *transmission attempt*, or simply *transmits*, at time t.

The interference between simultaneous transmissions has the following structure. If both vertices i and j transmit at time t, then the j's transmission is *erased* by i's transmission with a fixed probability ϕ_{ij}, $0 \le \phi_{ij} \le 1$. (By convention, $\phi_{ii} = 0$.) If vertex i is inactive in slot t, it does not interfere in any way with other nodes' transmissions in that slot. A transmission of vertex j at time t *succeeds*, if it is not erased by a simultaneous transmission of any other vertex, and it *fails* otherwise. It will be convenient to consider the *interference graph* $(\mathcal{I}, \mathcal{E})$, where directed *edge* (i, j) belongs to $\mathcal{E}$ if and only if $\phi_{ij} > 0$, that is when transmissions by i may erase those by j. For every $i \in \mathcal{I}$, we denote by

$$\mathcal{I}_i \doteq \{j \in \mathcal{I} \mid \phi_{ij} > 0\} \equiv \{j \in \mathcal{I} \mid (i, j) \in \mathcal{E}\}$$

the set of those vertices which may be affected by (the transmissions of) vertex i.

IV. RANDOM-ACCESS STRATEGY AND ITS THROUGHPUT REGION

Consider the following distributed, purely random access, "Slotted Aloha-type" strategy, parameterized by the vector of access probabilities

$$p = (p_1, \ldots, p_I) \in [0,1]^I. \tag{2}$$

The strategy is such that each vertex i in each time slot transmits with probability p_i, independently of other vertices and of the past history. (One can see that, under this strategy, our model is a generalization of the standard Slotted Aloha model, the latter being a special case with $\phi_{ij} = 1$ for all i and j.) The average *throughput* of vertex i, i.e., the average rate of successful transmissions, is the following function of p:

$$\mu_i(p) = p_i \prod_{j \neq i} (1 - \phi_{ji} p_j) \equiv p_i \prod_{j:\ i \in \mathcal{I}_j} (1 - \phi_{ji} p_j). \tag{3}$$

We denote the throughput vector corresponding to $p \in [0,1]^I$ by $\mu(p) \doteq (\mu_1(p), \ldots, \mu_I(p))$. Clearly, the function $\mu = \mu(p)$ is continuous for $p \in [0,1]^I$.

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

$$M \doteq \{\mu' \in [0,1]^I \mid \mu' \leq \mu(p) \text{ for some } p \in [0,1]^I\}. \tag{4}$$

We denote by

$$M^* \doteq \{\mu^* \in M \mid \mu^* \leq \mu' \in M \text{ implies } \mu' = \mu^*\} \tag{5}$$

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 paper. Also of interest is the larger, *outer* boundary M^{**} of M defined as

$$M^{**} \doteq \{\mu^* \in M \mid \mu^* \leq \mu' \in M \text{ implies } \mu'_i = \mu^*_i \text{ for some } i\}. \tag{6}$$

Obviously, $M^* \subseteq M^{**} \subseteq M$. Finally, we denote by M^*_{++} and M^{**}_{++} the subsets of M^* and M^{**}, respectively, consisting of vectors with all strictly positive components:

$$M^*_{++} \doteq M^* \cap R^I_{++}, \quad M^{**}_{++} \doteq M^{**} \cap R^I_{++}. \tag{7}$$

The following proposition describes some basic properties of and relations between sets M, M^* and M^{**}.

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

(ii) Set M^* is non-empty. For any $\mu^* \in M^*$ there exists $p \in [0,1]^I$ such that $\mu^* = \mu(p)$.

(iii) We have $M^* \subseteq \overline{M^*_{++}} \subseteq \overline{M^{**}_{++}} \subseteq M^{**}$.

Proof. (i) Since function $\mu = \mu(p)$ is continuous, the image of $[0,1]^I$ under mapping $\mu(\cdot)$ is a compact set. This and the definition of M easily imply compactness of M.

(ii) The fact that M^* contains at least one element follows from the compactness of M. The second statement follows directly from the definition of M.

(iii) Let us prove $M^* \subseteq \overline{M^*_{++}}$. For a fixed $\mu^* \in M^*$ let us pick p^* such that $\mu^* = \mu(p^*)$. Let us choose arbitrary sequence $p^{(\ell)}, \ell = 1, 2, \ldots$ such that $p^{(\ell)} \to p^*$ (and then $\mu(p^{(\ell)}) \to \mu^*$) and $0 < p_i^{(\ell)} < 1$ for all i and ℓ. For each ℓ and i, $\mu_i(p^{(\ell)}) > 0$. Let us choose a sequence $p^{(\ell),*}, \ell = 1, 2, \ldots$ such that $\mu(p^{(\ell),*}) \geq \mu(p^{(\ell)})$ and $\mu(p^{(\ell),*}) \in M^*_{++}$ for all l. It is easily seen that $\mu(p^{(\ell),*}) \to \mu^*$, which proves the first inclusion.

The second inclusion of (iii) is trivial. To prove $\overline{M^{**}_{++}} \subseteq M^{**}$, consider a converging sequence $\mu^{(\ell)} \to \mu^{**} \in M$, with $\ell = 1, 2, \ldots$ and $\mu^{(\ell)} \in M^{**}_{++}$ for all ℓ. Suppose $\mu^{**} \notin M^{**}$. Then, $\mu^{**} < \mu'$ for some $\mu' \in M$. Then, $\mu^{(\ell)} < \mu'$ for all large ℓ. The latter contradicts the fact that $\mu^{(\ell)} \in M^{**}_{++}$. The proof of the last inclusion of (iii) is complete. ∎

V. Non-convexity (in general) of M and $R_+^I \setminus M$.

As mentioned earlier, our model is a generalization of the classical Slotted Aloha random access model (cf. [11]), which is a special ("full interference") case with $\phi_{ij} = 1$ for all i and j. The throughput region M for the Slotted Aloha model has been characterized in [11], and it is non-convex. Therefore, *region M can be non-convex for our general model*. (In some special cases M can be convex. For example, in the extreme, "zero-interference" case, when $\phi_{ij} = 0$ for all i and j, M is simply the I-dimensional cube.) It is also known that, for the Slotted Aloha model, the complement of M in the positive orthant, $R_+^I \setminus M$ is convex. (This fact was conjectured in [11], and then proved in [12].) It is easy to see that, *for our general model, region $R_+^I \setminus M$ can be non-convex as well*.

Indeed, consider the system shown in Figure 1(a) consisting of 3 interfering links, with the corresponding interference graph shown in Figure 1(b). Further, assume that the interference is "0-1" type, i.e., ϕ_{21}, ϕ_{23} and ϕ_{32} are 1 and the rest 0. The corresponding throughput region M is plotted in Figure 2, which indicates that both M and $R_+^3 \setminus M$ for this system are non-convex (as evident from the faces $\mu_1 = 0$ and $\mu_2 = 0$, respectively). Furthermore, it can be seen that M^{**} is strictly bigger than M^* (for instance, $\{\mu_1 = 1, \mu_2 = 0, \mu_3 \in [0, 1)\} \subset M^{**} \setminus M^*$).

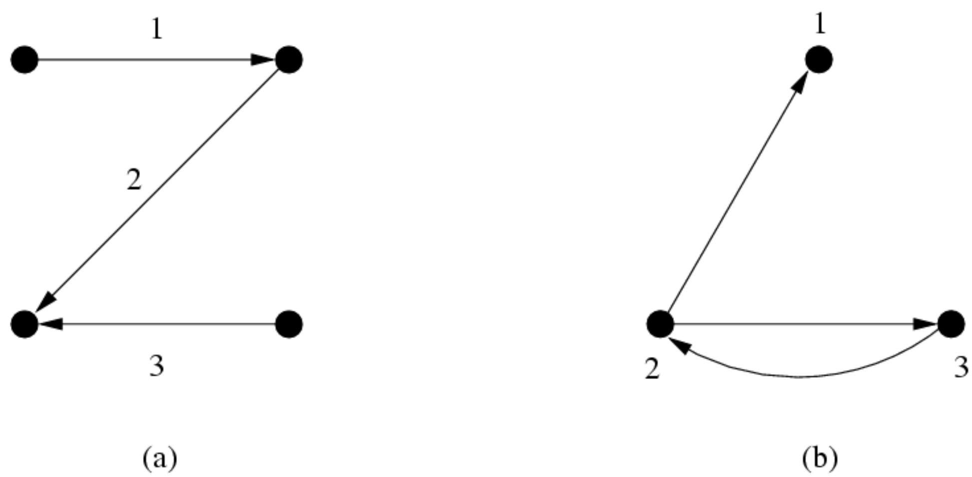

Fig. 1. (a) A 3-link system (b) the corresponding interference graph.

The non-convexity of region M necessitates a weaker notion of weighted proportional fairness, as discussed at the end of Section VI.

VI. Optimal Solution for the Weighted Proportional Fair Objective

Simple facts, summarized below in Theorem 2, serve as starting point for the characterization of Pareto boundary M^*. These facts basically follow from the observation that, given expression (3), the problem of choosing access probabilities so as to maximize $\sum w_i \log \mu_i$, with arbitrary fixed weights w_i, is quite tractable. (This optimization objective is sometimes called *weighted proportional fairness*. See remark at the end of this Section VI.)

Theorem 2: (i) Let a vector $w = (w_1, \ldots, w_2) \in R_{++}^I$ be fixed. Consider a function

$$F(\mu) \doteq \sum_{i \in \mathcal{I}} w_i \log \mu_i, \quad \mu \in R_+^N. \tag{8}$$

Then, the following holds.
(i.a) There exists a unique vector $p(w) = (p_1(w), \ldots, p_I(w))$ solving the optimization problem

$$\max_{p \in [0,1]^I} F(\mu(p)). \tag{9}$$

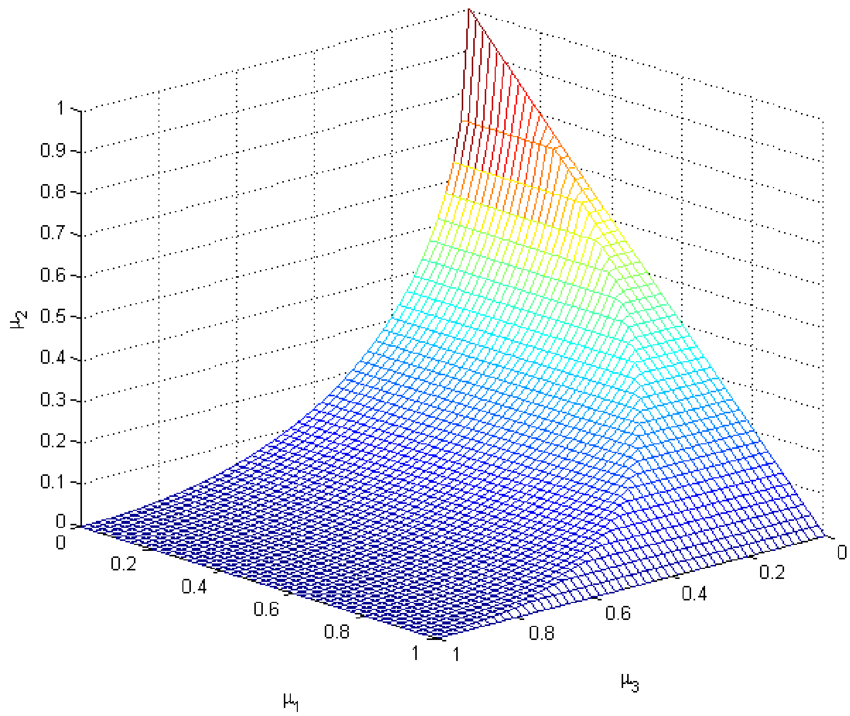

Fig. 2. The throughput region M of the 3-link system in Figure 1.

This vector $p(w)$ is as follows. Each component $p_i(w)$ is the unique solution of the one-dimensional maximization problem

$$\max_{p_i \in [0,1]} f_i(p_i), \tag{10}$$

where

$$f_i(p_i) \doteq w_i \log p_i + \sum_{j \in \mathcal{I}_i} w_j \log(1 - \phi_{ij} p_i). \tag{11}$$

More explicitly, $p_i(w)$ is equal to the unique solution of the equation

$$f_i'(p_i) = \frac{w_i}{p_i} - \sum_{j \in \mathcal{I}_i} \frac{w_j \phi_{ij}}{1 - \phi_{ij} p_i} = 0, \tag{12}$$

if $f_i'(1-) < 0$, and $p_i(w) = 1$ otherwise. In particular, for any $i \in \mathcal{I}$, we always have $p_i(w) > 0$ and, in addition, if $\phi_{ij} = 1$ for at least one $j \neq i$, then $p_i(w) < 1$.
(i.b) Vector $\mu(p(w)) \in M^*_{++}$.

(ii) Suppose we are in the conditions of (i) and, additionally, each parameter ϕ_{ij} is either 1 or 0. Then,

$$p_i(w) = \frac{w_i}{w_i + \sum_{j \in \mathcal{I}_i} w_j}. \tag{13}$$

(iii) Function $p(w),\ w \in R^I_{++}$, where $p(w)$ is defined in (i), is continuous. (Consequently, function $\mu(p(w)),\ w \in R^I_{++}$, is continuous.) In addition, $p(w)$ is invariant with respect to scaling by a positive constant, namely

$$p(cw) = p(w) \quad \forall c > 0. \tag{14}$$

Proof. (i.a) Using (3), we can write

$$F(\mu(p)) = \sum_{i\in\mathcal{I}} w_i \left[\log p_i + \sum_{j:\ i\in\mathcal{I}_j} \log(1-\phi_{ji}p_j) \right]$$

or, grouping for each i all the terms containing p_i,

$$F(\mu(p)) = \sum_{i\in\mathcal{I}} f_i(p_i), \tag{15}$$

where $f_i(p_i)$ is defined in (11). We see that the optimization problem (9) indeed decomposes into one-dimensional problems (10), separately for each p_i. Solution $p_i(w)$ to each problem (10) is unique, because the maximized function $f_i(p_i)$ is strictly concave continuously differentiable in $(0,1)$ (with the derivative $f_i'(p_i)$ as in (12)), $f_i(p_i)$ is upper bounded in $(0,1)$, $f_i'(0+) = +\infty$ and, finally, $f_i'(1-)$ is either finite or equal to $-\infty$. The (almost) explicit characterization of $p_i(w_i)$ via the derivative $f_i'(p_i)$, follows from the above listed properties of $f_i(\cdot)$. The remaining statement easily follows from this characterization.

(i.b) Denote $\mu^* \doteq \mu(p(w))$. It follows from the last property in (i.a), that $\mu^* \in R^I_{++}$. (This also follows from the fact that $F(\mu) = -\infty$ if $\mu_i = 0$ for at least one i.) We also must have $\mu^* \in M^*$, because otherwise, by the definition of M, we could find $p' \in [0,1]^I$ such that $\mu(p') \geq \mu^*$ and strict inequality $\mu_i(p') > \mu_i^*$ holds for at least one i. This would contradict to the fact that $p(w)$ maximizes $F(\mu(p))$, because $F(\mu)$ is strictly increasing in each μ_i. Thus, $\mu^* \in M^*_{++}$.

(ii) This is a corollary of (i.a). The case $\mathcal{I}_i = \emptyset$ is trivial. If $\mathcal{I}_i$ is non-empty, we have $f_i'(1-) = -\infty$, and can solve (12) explicitly to obtain (13).

(iii) It is easy to see that the dependence $f_i'(\cdot) = f_i'(\cdot; w)$ of function f_i' on $w \in R^I_{++}$ is such that the following property holds. Suppose we have a sequence $w^{(\ell)} \to w \in R^I_{++}$. Then, the convergence of functions $f_i'(\cdot; w^{(\ell)})$ to $f_i'(\cdot; w)$ is uniform in a neighborhood of $p(w)$. Because all functions $f_i'(\cdot; w^{(\ell)})$ and $f_i'(\cdot; w)$ are strictly decreasing, we must have $p(w^{(\ell)}) \to p(w)$. The scaling property (14) follows trivially from (i.a), because scaling of w simply scales each function $f_i(\cdot)$. ∎

Remark on the Notion of Weighted Proportional Fairness. A throughput vector $\mu^* \in M$, solving optimization problem (21) (in Theorem 2) for a given weight vector $w \in R^I_{++}$, that is

$$\mu^* = \arg\max_{\mu\in M} \sum_i w_i \log \mu_i, \tag{16}$$

is often called a *weighted proportional fair* throughput vector. Sometimes, however, the notion of a weighted proportional fair throughput vector μ^* is defined differently, as follows (cf. [10]):

$$\mu^* \text{ is such that } \sum_i w_i \frac{\mu_i - \mu_i^*}{\mu_i^*} \leq 0,\ \forall \mu \in M. \tag{17}$$

We would like to point out that, when M is *non-convex*, which is a typical situation for our general model (see Remark in Section IV), *the definitions (16) and (17) are not equivalent*. Namely, the following proposition holds. (We omit a rather straightforward proof.)

Proposition 3: Suppose, $M \in R^I_+$, $M \cap R^I_{++} \neq \emptyset$, and M is a compact set. Let $w \in R^I_{++}$ be fixed. Then, for a point $\mu^* \in M$ we have:

(i) Condition (17) implies (16).

(ii) If M is convex, condition (16) implies (17).

(iii) There exists $\mu^* \in M$ satisfying (16).

Thus, if M is non-convex, condition (17) is stronger in general than (16). In fact, in many cases of interest, as for example in the case of Slotted Aloha model ("full interference", all $\phi_{ij} = 1$), the points μ^* satisfying the stronger condition (17) do not exist.

VII. Characterization of the Pareto Boundary M^*

According to Theorem 2(i.b), all vectors of the form $\mu(p(w))$ with $w \in R^I_{++}$ lie on the Pareto boundary M^* (in fact, in the subset M^*_{++}) of the throughput region M. The natural question is: *Do vectors* $\mu(p(w))$ *in fact "fill" the entire set* M^*_{++}*?* If the answer is yes, this would be very useful, because it would mean that just by changing parameters w_i's, one can "steer" the vertex throughputs to a desired point on the Pareto boundary of the throughput region. As we show below, the answer is in fact *yes* for many cases of interest.

To this end, it is useful to consider log-throughput region, namely

$$U \doteq \{\log x \mid x \in M \cap R^I_{++}\}, \tag{18}$$

where log is understood component-wise. The Pareto boundary of U is

$$U^* \doteq \{\log x \mid x \in M^*_{++}\}. \tag{19}$$

Lemma 4: 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 $i \in \mathcal{I}$, $\log \mu_i(p)$ is a concave scalar function of the vector 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. ∎

Since $p(w)$ is invariant with respect to scaling of w (see Theorem 2(iii)), in what follows we usually restrict the domain R^I_{++} of $p(w)$ to the normalized set

$$B \doteq \{w \in R^I_{++} \mid \sum_i w_i = 1\}. \tag{20}$$

Using graph-theoretic terminology, we say that vertex j is *reachable* from vertex i, and write this as $(i \rightarrow j)$, if there exists a path from i to j along the (directed) edges of the interference graph $(\mathcal{I}, \mathcal{E})$. By convention, $(i \rightarrow i)$ for all i. Graph $(\mathcal{I}, \mathcal{E})$ is called *strongly connected* if all vertices are reachable from each other, that is, $(i \rightarrow j)$ for all $i, j \in \mathcal{I}$.

Remark. To avoid confusion, we want to emphasize that strong connectedness of interference graph $(\mathcal{I}, \mathcal{E})$ does *not* mean that "vertices can send or pass messages to each other." (In our model, vertices model different communication links of a network.) Presence of an edge $(i, j) \in \mathcal{E}$ means that vertex i causes interference to vertex j, that is, a transmission in i may erase a simultaneous transmission in j.

Lemma 5: Suppose the interference graph $(\mathcal{I}, \mathcal{E})$ is strongly connected. Suppose p is such that the corresponding $\mu^* = \mu(p) \in M \cap R^I_{++}$. (This in particular guarantees that all $p_i > 0$.) Then for any $k \in \mathcal{I}$ there exists p' such that $\mu' = \mu(p')$ has $\mu'_k < \mu^*_k$ and $\mu'_i > \mu^*_i$ for all $i \neq k$.

Proof. Let us initially set $p' = p$ and then change it in steps as follows. First, slightly decrease p'_k. It is easy to see that $\mu' = \mu(p')$ becomes such that $\mu'_k < \mu^*_k$, $\mu'_i > \mu^*_i$ for each $i \in \mathcal{I}_k$, and $\mu'_i = \mu^*_i$ for the remaining i. Next, we pick a vertex $j \in \mathcal{I}_k$ such that $\mathcal{I}_j$ contains $i \notin \mathcal{I}_k \cup \{k\}$; such j exists by strong connectedness. We will decrease p'_j by a small enough amount such that conditions $\mu'_k < \mu^*_k$ and $\mu'_j > \mu^*_j$ are preserved; μ'_i will increase for all $i \in \mathcal{I}_j$. Thus, we increased the subset of vertices $i \neq k$ with $\mu'_i > \mu^*_i$. We can continue (using strong connectedness) until $\mu'_i > \mu^*_i$ for all $i \neq k$, while preserving condition $\mu'_k < \mu^*_k$. ∎

Theorem 6: Suppose the interference graph $(\mathcal{I}, \mathcal{E})$ is strongly connected. Then, the following holds.
(i) $M^*_{++} = M^{**}_{++}$.
(ii) M^*_{++} is the image of B under the (continuous) mapping $\mu(p(w))$. In other words, for any $\mu^* \in M^*_{++}$ there exists $w \in B$ such that $\mu^* = \mu(p(w))$. (There is no claim of uniqueness of w.)

(iii) If $\mu^* \in M^*_{++}$ and $w \in B$ are such that $\mu^* = \mu(p(w))$, then $p(w)$ is the unique solution of equation $\mu^* = \mu(p)$, and μ^* is the unique solution of the optimization problem

$$\max_{\mu \in M} F(\mu). \tag{21}$$

Proof. (i) We need to show that $\mu^{**} \in M^{**}_{++}$ implies $\mu^{**} \in M^*_{++}$. Suppose not. Then, for some $\mu^* \in M^*_{++}$, we have $\mu^{**} \leq \mu^*$ and $\mu^{**}_k < \mu^*_k$ for at least one k. Using Lemma 5, we can construct vector $\mu' \in M$ such that $\mu^{**} < \mu'$, which contradicts to condition $\mu^{**} \in M^{**}_{++}$.

(ii) If $\mu^* \in M^*_{++}$, then $u^* = \log \mu^* \in U$. Since U is convex and all $u \leq u^*$ are in U, there exist vector w with non-negative components such that u^* solves problem

$$\max_{u \in U} w \cdot u. \tag{22}$$

In fact, all components of w must be positive. Otherwise, if some $w_k = 0$, then using Lemma 5, we could construct vector $\mu' \in M$ such that $\mu'_k < \mu^*_k$ and $\mu'_i > \mu^*_i$ for all $i \neq k$, and then $u' = \log \mu'$ would be such that $w \cdot u' > w \cdot u^*$, a contradiction. Therefore, without loss of generality, $w \in B$. But then, by Theorem 2(i.a) we must have $\mu^* = \mu(p(w))$.

(iii) The solution to $\mu^* = \mu(p)$ maximizes $\sum w_i \log \mu_i(p)$ over all p, and therefore, by Theorem 2(i.a), it is unique and is given by $p(w)$. Similarly, μ^* is the unique solution of (9), because it is must be equal to $\mu(p)$ with $p = p(w)$. ∎

Remark on the Strong Connectedness Assumption. The analysis of this section can be expanded to produce some results for the case when the interference graph is not strongly connected. Formally, in this case $w \in B$ such that $\mu^* = \mu(p(w))$ might *not* exist for some $\mu^* \in M^*_{++}$. This property, however, still holds "in a limit sense." Namely, it is not hard to show that for any $\mu^* \in M^*_{++}$ there exists a *sequence* of $w^{(\ell)} \in B,\ \ell = 1, 2, \ldots$, such that $\mu(p(w^{(\ell)})) \to \mu^*$.

We illustrate this fact using the example given in Section V. The interference graph is not strongly connected. Consider the following throughput vector $\mu^* = (1/2, 1/4, 1/4) = \mu(p^*)$, where $p^* = (1, 1/2, 1/2)$. It is easy to see that $\mu^* \in M^*_{++}$. However, this throughput allocation cannot be achieved as $\mu(p(w))$ with any $w \in R^I_{++}$. On the other hand, $\mu^* = \lim_{\ell \to \infty} \mu(p(w^{(\ell)}))$ where $w^{(\ell)} = (1, \ell, \ell)$. The intuition is that μ^* is obtained by optimizing $\log(\mu_2) + \log(\mu_3)$, while "ignoring" vertex 1; then we maximize μ_1. This is equivalent to optimizing $\log(\mu_1) + w_2 \log(\mu_2) + w_3 \log(\mu_3)$ with w_2 and w_3 being equal and infinitely large.

We also note that, obviously, even if the interference graph is strongly connected, both M and $R^I_+ \setminus M$ can be non-convex. For instance, in the example in Section V, if we make $\phi_{13} = \delta > 0$, with δ very small, this will make the interference graph strongly connected, while keeping both M and $R^I_+ \setminus M$ non-convex.

VIII. Smoothness Properties of M^*

Theorem 7: Suppose the interference graph $(\mathcal{I}, \mathcal{E})$ is strongly connected. Let $w^* \in B$ is such that $p^* = p(w^*) \in (0,1)^I$. (Recall that all $p^*_i > 0$ "automatically.") Also, denote $\mu^* = \mu(p^*) = \mu(p(w^*))$. Then, the following properties hold.

(i) Vector w^* is the unique solution (within B) of equations $p^* = p(w)$ and $\mu^* = \mu(p(w))$.

(ii) We have

$$w^* J = 0, \tag{23}$$

where the $I \times I$ matrix $J = (J_{ij},\ i \in \mathcal{I},\ j \in \mathcal{I})$ is defined as follows:

$$J_{ij} = \begin{cases} 1/p^*_i, & \text{if } i = j \\ -\phi_{ji}/(1 - \phi_{ji} p^*_j), & \text{if } i \in \mathcal{I}_j \\ 0, & \text{otherwise} \end{cases} \tag{24}$$

Moreover, w^* is the unique (up to scaling) left eigenvector of matrix J and, consequently,

$$rank(J) = I - 1.$$

(iii) The Jacobi matrix

$$(\partial\mu/\partial w) \doteq (\partial\mu_i/\partial w_j,\ i \in \mathcal{I},\ j \in \mathcal{I})$$

of the function $\mu(p(w))$ at point w^* has the following form:

$$(\partial\mu/\partial w) = (\partial\mu/\partial p)(\partial p/\partial w), \tag{25}$$

where

$$(\partial\mu/\partial p) = diag(\mu^*)J, \tag{26}$$

$$(\partial p/\partial w) = diag(b)J^T, \tag{27}$$

$b = (b_1, \dots, b_I)$,

$$b_i = \frac{w_i^*}{[p_i^*]^2} + \sum_{j\in\mathcal{I}_i} \frac{w_j^*\phi_{ij}^2}{[1-\phi_{ij}p_i^*]^2} > 0. \tag{28}$$

In addition,

$$rank(\partial\mu/\partial w) = I - 1 \tag{29}$$

and

$$(w^*/\mu^*)(\partial\mu/\partial w) = 0. \tag{30}$$

(In other words, the values of the differential $(d\mu)^T = (\partial\mu/\partial w)(dw)^T$ "fill" the $(I-1)$-dimensional subspace, orthogonal to vector (w^*/μ^*).)

Proof. (i) Since all components of p^* are strictly between 0 and 1, for any $w \in B$ not equal to w^* we must have $p(w) \neq p(w^*) = p^*$. Indeed, consider arbitrary $w \in B$ not equal to w^*. Denote by $\mathcal{I}^{(1)}$ the subset of those $i \in \mathcal{I}$, for which the ratio w_i/w_i^* is maximal. Clearly, $\mathcal{I}^{(1)}$ is a strict subset of $\mathcal{I}$. Then, because $(\mathcal{I},\mathcal{E})$ is strongly connected, there exists an edge $(i,j) \in \mathcal{E}$ such that $i \in \mathcal{I}^{(1)}$ and $j \in \mathcal{I}\setminus\mathcal{I}^{(1)}$. Then, it is easily seen from (11), that $p_i(w) > p_i(w^*)$. Thus, solution w^* of $p^* = p(w)$ is the unique. Solution to $\mu^* = \mu(p(w))$ is also unique, because by Theorem 6(iii) $p(w)$ must be the unique solution of $\mu^* = \mu(p)$, implying $p(w) = p^*$ and then $w = w^*$.

(ii) Equation (23) follows from the definition of vector p^*. Indeed, i-th component of row-vector w^*J is equal to

$$\frac{w_i^*}{p_i^*} - \sum_{j\in\mathcal{I}_i} \frac{w_j^*\phi_{ij}}{1-\phi_{ij}p_i^*} = 0, \tag{31}$$

by (12).

Now, we show that w^* is the unique (up to scaling) non-zero vector satisfying (23). Suppose not, and w' is another such vector. If we choose $\epsilon > 0$ sufficiently small, then vector $w'' = w^* + \epsilon w' \in R_{++}^I$ and it is *not* a scaled version of w^*. Then, we have $w''J = 0$, which implies, for each i, (31) with w^* replaced by w''. This, however, means $p^* = p(w'')$, which contradicts statement (i).

(iii) The form of the Jacobi matrices in (26) and (27) is verified directly from (3) and (12). The rank of $(\partial\mu/\partial w) = diag(\mu^*)\ J\ diag(b)\ J^T$ is equal to the rank of the matrix $[J\ diag(\sqrt{b})][J\ diag(\sqrt{b})]^T$. The latter matrix has the same set of left eigenvectors as the matrix $[Jdiag(\sqrt{b})]$, and therefore its rank is $I-1$, which proves (29). Equation (30) holds because

$$(w^*/\mu^*)diag(\mu^*)J = w^*J = 0.$$

∎

Corollary 8: Suppose, we are in the conditions of Theorem 7. Then, "in a small neighborhood of point μ^*, the boundary M_{++}^* is a smooth $(I-1)$-dimensional surface, parameterized by the elements $w \in B$ close to w^*." The formal statement is as follows. For a given $\beta > 0$, let us denote

$$B(\beta) \doteq \{w \in B \mid \|w - w^*\| < \beta\}.$$

Then, if $\beta > 0$ is small enough, the following properties hold.

(i) The set $M^*_{++}(\beta) \subseteq M^*_{++}$, defined as the image of $B(\beta)$ under the mapping $\mu(p(w))$, is a smooth $(I-1)$-dimensional surface in R^I_{++}.
(ii) Mapping $\mu(p(w))$ defines a homeomorphism (mutually continuous one-to-one mapping) between $B(\beta)$ and $M^*_{++}(\beta)$.
(iii) For any $\mu' \in M^*_{++}(\beta)$, the hyperplane

$$\mu \cdot (w'/\mu') = 1$$

is the (unique) tangent hyperplane to the surface $M^*_{++}(\beta)$ at point μ'.
(iv) For a sufficiently small $\delta > 0$, we have

$$M^*_{++} \cap \{\|\mu - \mu^*\| < \delta\} = M^*_{++}(\beta) \cap \{\|\mu - \mu^*\| < \delta\}.$$

Proof. We only need to observe that if $\beta > 0$ is small enough, then, by continuity of functions $p(w)$ and $\mu(p(w))$, every point $w \in B(\beta)$ is within the conditions of Theorem 7. The dependence of the Jacobian, and then that of the orientation of the $(I-1)$-dimensional differential $(d\mu)^T$, on w (at $\mu(p(w))$) is also continuous. Finally, all points $\mu' \in M^*_{++}$ sufficiently close to μ^* must have the corresponding w' close to w^*, because otherwise we could get a contradiction to the fact that w^* is the unique solution of $\mu^* = \mu(p(w))$. ∎

Corollary 9: Suppose the interference graph $(\mathcal{I}, \mathcal{E})$ is strongly connected and, in addition, each parameter ϕ_{ij} is either 1 or 0. Then, the entire boundary M^*_{++} is a smooth $(I-1)$-dimensional surface. Function $\mu(p(w))$ is differentiable for all $w \in B$ and defines a homeomorthism between B and M^*_{++}. For any pair of $w^* \in B$ and $\mu^* \in M^*_{++}$ corresponding to each other (that is, $\mu^* = \mu(p(w^*))$), all properties described in Theorem 7 and Corollary 8 hold.

Proof. Assumptions of the theorem, along with Theorem 2(ii), imply that for any $w \in B$, $p_i(w) \in (0,1)$. This puts each $w^* \in B$ within the assumptions of Theorem 7 and Corollary 8. ∎

References

[1] V. Anantharam. The stability region of the finite-user slotted Aloha protocol. *IEEE Trans. Inform. Theory,* Vol. 37, (1991), No. 3, pp. 535–540.
[2] N. Abramson. The ALOHA system – Another alternative for computer communications. *Proc. AFIPS Conf.*, Fall Joint Computer Conference, Vol. 37, (1970), pp. 281-285.
[3] S. Ghez, S. Verd, S. Schwartz. Stability Properties of Slotted-Aloha with Multipacket Reception Capability. *IEEE Trans. Autom. Control*, vol. 33, no. 7, p. 640-649, July 1988.
[4] 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.
[5] P. Gupta, A. L. Stolyar. Optimal Throughput Allocation in General Random-Access Networks. *Proceeding of CISS'2006*, Princeton, March 2006.
[6] J. Liu, A. L. Stolyar, M. Chiang, H. V. Poor. Queue Back-Pressure Random Access in Multi-Hop Wireless Networks: Optimality and Stability. *IEEE Trans. Inf. Theory*, 2009, Vol.55, No.9, pp. 4087-4098.
[7] L. Jiang and J. Walrand. A distributed CSMA algorithm for throughput and utility maximization in wireless networks. *Proceedings of the 46th Annual Allerton Conference*, Urbana-Champaign, 2008.
[8] L. Jiang, D. Shah, J. Shin, J. Walrand. Distributed random access algorithm: scheduling and congestion control. *IEEE Trans. Inform. Theory,* Vol. IT-56, (2010), no. 12, pp. 6182–6207.
[9] 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.
[10] 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.
[11] J. Massey and P. Mathys. The collision channel without feedback. *IEEE Trans. Inform. Theory,* Vol. IT-31, (1985), no. 2, pp. 192–204.
[12] 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.
[13] A. L. Stolyar. Dynamic Distributed Scheduling in Random Access Networks. *Journal of Applied Probability*, 2008, Vol.45, No.2, pp. 297-313.
[14] 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.

PLACE
PHOTO
HERE

Piyush Gupta received the Bachelor of Technology degree in electrical engineering from the Indian Institute of Technology, Bombay, in 1993, the Master of Science degree in computer science and automation from the Indian Institute of Science, Bangalore, in 1996, and the Ph.D. degree in electrical and computer engineering from the University of Illinois, Urbana-Champaign, in 2000.

From 1993 to 1994, he was a design engineer at the Center for Development of Telematics, Bangalore. Since September 2000, he has been a Member of Technical Staff in the Mathematics of Networks and Communications Research Department at Alcatel-Lucent, Bell Laboratories, Murray Hill, NJ. His research interests include wireless networks, network information theory, and learning and adaptive systems.

PLACE
PHOTO
HERE

Alexander L. 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 and service systems.