Asymptotic Properties of Proportional-Fair Sharing Algorithms: Extensions of the Algorithm

Harold J. Kushner*
Division of Applied Mathematics
Brown University
Providence RI 02912
hjk@dam.brown.edu

Philip A. Whiting
Bell Labs, Lucent Technologies
600 Mountain Ave, Room 2C-317
Murray Hill NJ 07974-0636
pawhiting@lucent.com

Abstract

We are concerned with the allocation of transmitter time and power for randomly time varying mobile data communications. Time is divided into small scheduling intervals, called slots, and information on the channel rates for the various users is available at the start of the slot, when the user selections are made. There is a conflict between selecting the user set that can get the most immediate data through and helping users with poor average rates. The Proportional Fair Sharing method (PFS) deals with such conflicts. In [5, 6] the convergence and basic qualitative properties were analyzed. Stochastic approximation results were used to analyze the long term properties. The paths of the (suitably interpolated) throughputs converge to the solution of an ODE, akin to a mean flow. The ODE has a unique equilibrium point. It is asymptotically stable and optimal in that it is the maximizer of a concave utility function. There is a large family of such algorithms, each member corresponding to a concave utility function. The basic idea of PFS extends to many systems of current importance for which it was not originally intended, and a variety of such extensions are treated here to illustrate the possibilities. One might have minimal throughput constraints, nonlinear dependence of rate on allocated power, minimal SNR requirements, etc. In some recent applications, the number of slots in a scheduling intervals is random, and the length is not known when the selection is made. The form of the PFS rule is adapted to the application. Then the basic results continue to hold. The asymptotic properties of the ODE characterize the behavior of the algorithm.

1 Introduction and Background

Consider the problem where there are N users wishing to transmit data from a single base station to N mobile destinations, and the rates of transmission are randomly varying. Time is divided into small scheduling intervals, called slots. Until further notice, in each slot one user is chosen to transmit. New systems estimate the rates by estimating the SNR, by use of a pilot signal, with a very short delay. Rayleigh fading rates of a few tens of Hertz can be accommodated. The selection of the user is based on a balance between the current possible rates and "fairness." The Qualcom algorithm, *proportional fair sharing* (PFS). performs this by comparing the given rate for each user with its average throughput to date.

*Partially supported by Army Research Office Contract DAAD-19-02-1-0425 and National Science Foundation Grants ECS 0097447.

The evolution of the throughput sequence under PFS can be represented in a recursive manner, as a stochastic approximation (SA) [7] algorithm, and results of SA used to analyze its properties. [5, 6] dealt rigorously with the asymptotic and qualitative properties. The sequence of suitably interpolated throughputs converges to the solution of a "mean" ODE, which has a unique and asymptotically stable equilibrium point $\bar{\theta}$. The throughput paths converge to $\bar{\theta}$. This in turn, gave other properties (e.g,, quantifying the scheduling gains over simple TDMA). The algorithm was optimal in that no other can achieve a larger long run value of an associated utility function than PFS. This paper continues the work in [5, 6]. The original algorithm assigns only one user at a time, at full power and with no constraints. But the way that fairness was defined and the interpretation in terms of a utility function can be exploited to obtain natural PFS algorithms for many other types of systems of current interest, with all of the analytical results carrying over. Section 2 reviews the past work, with special emphasis on a property of the system known as *cooperativeness.* Sections 3 and 4 concern a selection of extensions, opening possibilities to a wider exploitation of the basic ideas: Many users might be scheduled simultaneously; channel rates might be nonlinear functions of the allocated power; there might be constraints. The appropriate form of the PFS rule is developed. There are applications where the number of slots in a scheduling interval is actually random. In all these cases, the basic theorems carry over under appropriate conditions.[1]

2 The PFS Algorithm: Outline of Past Results

The original PFS algorithm. This section describes the original PFS algorithm and some past results. There are N users, each with infinite backlog,[2] sending data to mobiles via a single transmitter, with the possible rates of each varying randomly. Users are scheduled one at a time. If user i is selected in slot n, then it transmits $r_{i,n}$ units of data, where $\{r_{i,n}, n < \infty\}$ is a random sequence. The end of slot n is called *time* n. One definition of throughput up to time n for user i is[3]

$$\theta_{i,n} = \sum_{l=1}^{n} r_{i,l} I_{i,l}/n, \quad I_{i,l} = 1 \text{ indicates user } i \text{ chosen in slot } l, \tag{1}$$

$$\theta_{i,n+1} = \theta_{i,n} + \epsilon_n \left[I_{i,n+1} r_{i,n+1} - \theta_{i,n} \right], \quad \epsilon_n = 1/(n+1). \tag{2}$$

An alternative definition of throughput discounts past values as

$$\theta^{\epsilon}_{i,n} = (1-\epsilon)^n \theta^{\epsilon}_{i,0} + \epsilon \sum_{l=1}^{n} (1-\epsilon)^{n-l} r_{i,l} I_{i,l}, \quad \epsilon > 0 \text{ small discount factor}, \tag{3}$$

where θ^{ϵ}_{n} is the discounted throughput at time n. In recursive form, (3) is

$$\theta^{\epsilon}_{i,n+1} = \theta^{\epsilon}_{i,n} + \epsilon \left[I_{i,n+1} r_{i,n+1} - \theta^{\epsilon}_{i,n} \right]. \tag{4}$$

(2) and (4) are of the SA form [7], and the results of SA were used for their analysis. The approach is able to handle quite general queues, arbitrary channel rate processes,

[1] Alternative points of view, which do not seem suitable for the general queueing problem, but which can be useful if the cooperativeness property doesn't hold, are in [1, 10]. The conditions are stronger: for example, [1] seems to exclude common cases where the channels are independent.

[2] The case where data arrives at random and is queued until transmitted is covered in [6].

[3] If desired, ϵ_n in (2) can go to zero more slowly than $1/n$, useful for iterate averaging [5].

and algorithm formulations. The basic algorithm and results are typical of a large family, indexed by some concave utility function, and which is actually maximized by the associated algorithm. The set of algorithms allow a continuum of tradeoffs between current rate and throughput, and the same advantage holds for the extensions to follow. In [6] both (2) and (4) were treated, differing mainly in notation. We work only with (4); identical results hold for (2).

Let $d_i > 0$, which can be as small as desired. The original PFS algorithm chooses the user at time n maximizing in

$$\arg\max\{r_{i,n+1}/[\theta^{\epsilon}_{i,n} + d_i], i \leq N\} \tag{5}$$

Definitions and outline of results for (4). Define $\theta^{\epsilon}_n = \{\theta^{\epsilon}_{i,n}, i \leq N\}$ and $R_n = \{r_{i,n}, i \leq N\}$. The usual SA large time analysis uses continuous time interpolations. Define $\theta^{\epsilon}(\cdot)$ (components $\theta^{\epsilon}_i(\cdot), i \leq N$) by $\theta^{\epsilon}(t) = \theta^{\epsilon}_n$ for $t \in [n\epsilon, n\epsilon+\epsilon)$. Let ξ_n denote the data $\{R_l : l \leq n\}$ to time n, and E_n the expectation conditioned on ξ_n. The assumptions are just examples of the possibilities. See [5, 6] for weaker assumptions. The last part of A3 is used to avoid degeneracy and is unrestrictive. It holds if when a component θ_i is small enough then there is a nonzero chance that user i will be chosen, no matter what the values of the other components of θ. A1–A3 hold under Rayleigh fading. The continuity in A1 holds when successive rates differ by some "randomness" so that small changes in θ will change only slightly the conditional probability that any user is selected. (A3) is used to show that the limit point is unique. (6) is weaker than the law of large numbers. We will abuse terminology by using θ_i and θ_j for the i-th and j-th components of the vector θ.

A1. *For each i, n, the functions $h_{i,n}(\theta, \xi_n) = E_n r_{i,n+1} I_{\{r_{i,n+1}/(d_i+\theta_i) \geq r_{j,n+1}/(d_j+\theta_j), j\neq i\}}$, are continuous. Let $\delta > 0$. Then in $\{\theta : \theta_i \geq \delta, i \leq N\}$, the continuity is uniform in n, ξ_n.*

A2. *$\{R_n, n < \infty\}$ is stationary. Define $\bar{h}_i(\theta) = E r_i I_{\{r_i/r_j \geq (d_i+\theta_i)/(d_j+\theta_j), j\neq i\}}$ (stationary expectation). $\bar{h}(\cdot)$ is continuous.* (In fact, A3 implies that $h(\cdot)$ is Lipschitz s [6].) *Also,*

$$\lim_{m,n\to\infty} \frac{1}{m} \sum_{l=n}^{n+m-1} \left[E_n r_{i,l+1} I_{\{r_{i,l+1}/r_{j,l+1} \geq (d_i+\theta_i)/(d_j+\theta_j), j\neq i\}} - \bar{h}_i(\theta)\right] = 0 \tag{6}$$

in probability. There are small positive δ, δ_1 such that $\bar{h}_i(\theta) \geq \delta_1$ for $\theta_i \leq \delta$, $i \leq N$.

A3. *R_n is defined on some bounded set and has a bounded density.*

The limiting process. The next theorem is a standard result in SA. It says that the limit points of (4), are contained in those of the ODE (7).

Theorem 1. ([7, Theorems 2.2 and 2.3, Section 8.2].) *Assume algorithm* ((4), (5)), A1, and A2. *Then for any bounded set of initial conditions, any subsequence of $\theta^{\epsilon}(q_{\epsilon}\epsilon + \cdot)$ has a further subsequence that converges weakly to the set of limit points of the solution of*

$$\dot{\theta}_i = \bar{h}_i(\theta) - \theta_i, \quad i \leq N. \tag{7}$$

Comment on a monotonicity property. See [5, 6]. The proofs of Theorems 2–4 use a monotonicity property together with the strict concavity of the log function, properties that are essential for the the extensions. For vectors $X, Y \in I\!R^N$, write $X \geq Y$ (resp., $X > Y$) if $X_i \geq Y_i$ for all i (resp., and also $X \neq Y$). The ODE (7) has an important property, called the K condition. A function $f(\cdot)$ is said to satisfy the K-condition if for any x, y, i, with $x \leq y$ and $x_i = y_i$, $f_i(x) \leq f_i(y)$. In our case, $f(\theta) = \bar{h}(\theta) - \theta$, and the condition holds. The condition says simply the following: If θ_i is increased, then the other users are not less likely to be chosen. Systems satisfying the K-condition are said to be *cooperative.* The condition implies the following important monotonicity theorem.

Theorem 2. [9, Proposition 1.1] *Let $f(\cdot)$ be Lipschitz continuous and assume the K-condition for $\dot{x} = f(x)$. If $x(0) \le y(0)$ (resp., $<$), then $x(t|x(0)) \le x(t|y(0))$ (resp., $<$), where $x(t|y)$ is the solution under initial condition y.*

The next result says that the throughputs converge to unique limiting values

Theorem 3. *Assume rule* ((4), (5)) and A1–A3. *The ODE* (7) *is cooperative and the limit point $\bar{\theta}$ is unique and asymptotically stable. $\theta^\epsilon(t+\cdot) \to \bar{\theta}$ as $\epsilon \to 0$ and $t \to \infty$.*

Maximizing a utility function. Theorem 3 shows that there is a unique asymptotically stable limit point $\bar{\theta}$ of the ODE and algorithm. The algorithm is also optimal, as follows. PFS is a type of "ascent" algorithm for the strictly concave utility function

$$U(\theta) = \sum_i \log(d_i + \theta_i). \tag{8}$$

By a first order Taylor expansion,

$$U(\theta^\epsilon_{n+1}) - U(\theta^\epsilon_n) = \epsilon \sum_i [r_{i,n+1} I_{i,n+1} - \theta^\epsilon_{i,n}]/[d_i + \theta_{i,n}] + O(\epsilon_n^2). \tag{9}$$

Since $\sum_i I_{i,n+1} = 1$, to maximize the first order term one must choose $I_{i,n+1}$ by (5). The forms of the extensions are motivated by the maximization of the first order term in such expansions. Theorem 4 asserts that the (5) maximizes $U(\cdot)$ [5].

Theorem 4.[4] *Under* A1–A3, (5) *maximizes* $\lim_n U(\theta_n)$.

An example: Rayleigh fading. Consider independent users with received signal power determined by stationary Rayleigh fading and with additive noise of constant power. Let the rates be proportional to the SNR, with mean rates $1/\beta_i, i \le N$. Then $\bar{h}(\cdot)$ can be evaluated for $d_i = 0$ [6]: $\bar{\theta}_i = G(N)/N \cdot 1/\beta_i$ where $G(N) = \sum_{j=1}^N 1/j$. This represents a gain over classical TDMA, where $\theta_i = 1/[N\beta_i]$. This was noted in [4], under stronger conditions. The fact that the limit throughput is proportional to the mean possible rate is a consequence of the facts that the stationary distribution for Rayleigh fading is exponential, the channels mutually independent, and the scheduling rule is (5). In [3], under the same conditions it is shown that all users get the same fraction of slots, asymptotically. The proofs of these properties depend on the uniqueness of the limit.

3 Extensions

3.1. Nonlinear Dependence of Rate on Power. The Model and PFS Rule. We now allow several users to be scheduled in each slot, with a total power $\bar{u}$. Consider a channel with white-Gaussian noise and signal to noise ratio S/N, with an unrestricted signalling alphabet. Then the channel capacity is $\log(1+S/N)$.[5] Let u_i denote the power assigned to user i at time n, $\alpha_{i,n+1}$ the power attenuation from transmitter to mobile i, and $N_{i,n+1}$ the receiver noise power. Define $b_{i,n+1} = \alpha_{i,n+1}/N_{i,n+1}$. Then the corresponding

[4]All of the above results can be extended to algorithms based on any strictly concave utility function, not just (8). This yields a family of algorithms that allow different tradeoffs between the values of the current rates and throughputs [5].

[5]One can suppose that there are many symbol intervals per code word, and many code words per slot, so that the capacity formulas can be used. This formula is for an unrestricted signalling alphabet. But, for small SNR values (say, ≤ 4 db), it can be well approximated even by binary alphabets [8].

channel capacity for user i in slot $n+1$ is $vr_{i,n+1}(u) = v \log(1 + b_{i,n+1}u_i)$, where v denotes the number of symbol intervals per slot.[6] Redefine θ_i to be the throughput divided by v. Then (9) becomes

$$U(\theta^\epsilon_{n+1}) - U(\theta^\epsilon_n) = \epsilon \sum_i [\log(1 + b_{i,n+1}u_i) - \theta^\epsilon_{i,n}]/[\theta^\epsilon_{i,n} + d_i] + O(\epsilon^2). \tag{10}$$

We can suppose that the SNR determined from the pilot signal is $\alpha_{i,n+1}/N_{i,n+1}$, that for unit power. Due to the strict concavity in the u_i, the power might be distributed among more than one user. Using the motivation for (5) in terms of the maximization of the first order increment in the log utility function in (9), a natural analog of (5) is

$$\max_u \sum_i [\log(1 + b_{i,n+1}u_i)]/[\theta_i + d_i], \quad \text{subject to } \sum_i u_i \le \bar{u}, \quad u_i \ge 0. \tag{11}$$

Let $\bar{u}(\theta, b) = \{\bar{u}_i(\theta, b), i \le N\}$ denote the maximizer. Then (4) is replaced by

$$\theta^\epsilon_{n+1} = \theta^\epsilon_n + \epsilon \left[\sum_i \log(1 + b_{i,n+1}\bar{u}_i(\theta^\epsilon_n, b_{n+1})) - \theta^\epsilon_n \right]. \tag{12}$$

Computing the maximizer $\bar{u}(\theta, b)$ is not difficult, and we now turn to this problem.

Finding the maximizer in (11). For simplicity drop the n and ϵ, and write $b_{i,n+1}$ as b_i, and θ_i for $\theta_i + d_i$. First, drop the constraint $u_i \ge 0$ and form the Lagrangian

$$L(\lambda_i, \mu, \theta) = \sum_i [\log(1 + b_i u_i)]/[\theta_i] + \mu \left[\bar{u} - \sum_i u_i \right],$$

where μ is the multiplier. If $b_i = 0$, then drop the ith summand. Thus, we can suppose that all $b_i > 0, \theta_i > 0$. Differentiating with respect to u_i yields $b_i/[(1+b_i u_i)\theta_i] - \mu = 0,\ i \le N$. Thus, we need only solve $b_1/[(1 + b_1 u_1)\theta_1] = b_i/[(1 + b_i u_i)\theta_i]$, $2 \le i \le N; \sum_i u_i = \bar{u}$, or, equivalently,

$$\theta_i u_i - \theta_1 u_1 = [b_i \theta_1 - b_1 \theta_i]/b_1 b_i = c_j, \quad 2 \le i \le N; \ \sum_i u_i = \bar{u}, \tag{13}$$

(13) has a unique Lipschitz continuous solution $\bar{u}(\theta, b) = \{\bar{u}_i(\theta, b), i \le N\}$. [7]

Suppose that some components of the solution are negative. The fact that each summand in (12) depends only on a single component, together with the concavity of the log function, implies that these negative components will be zero for the problem where the nonnegativity constraint is imposed. So drop these components and resolve (13). Continue until all components are nonnegative. We can allow the b_i to be discrete-valued; they need not have a density. The density in (A3) was used only to assure the Lipschitz property of $\bar{h}(\cdot)$ dynamics. Here that is assured by the smoothness of $\bar{u}(\theta, b)$.

The limit ODE and convergence theorems. A1–A3 are not needed, except for the last part of A2. (6) is replaced by the following. Let b_n denote $\{b_{i,n}, i \le N\}$.

---

[6] We have not included the effects of mutual interference among the users. With orthogonal spreading sequences, the interference will be small. Also, in many cases, it does not depend on how the power is distributed, only on the total power. Then SIR replaces SNR and the results continue to hold. General interference can be handled in that the theorems remain valid. But the computation of the power allocation can be quite hard.

[7] If $U(\cdot)$ is replaced by any strictly concave utility function, then the resulting system still has the form of (13), but with different coefficients.

A4. *For any bounded and continuous real-valued function $f(\cdot)$, there is a constant $\bar{f}$:*

$$\lim_{n,m\to\infty} \frac{1}{m} \sum_{l=n}^{n+m-1} E_n \left[f(b_l) - \bar{f} \right] \to 0, \qquad \text{in probability.} \tag{14}$$

The limit ODE is

$$\dot{\theta}_i = \bar{h}_i(\theta) - \theta_i, \quad i \le N, \tag{15}$$

where $\bar{h}_i(\theta)$ is defined by (14) for $f(b_n) = \log(1 + b_{i,n}\bar{u}_i(\theta, b_n))$. The cooperative property and Theorems 1, 3, and 4 holds. The proof that (15) is the limit ODE is as in [5, 6], as is the proof that there is a unique equilibrium point $\bar{\theta}$ such that the paths starting at small θ all go to $\bar{\theta}$. Similarly, for the proof that all paths end up in $Q(\bar{\theta}) = \{\theta : \theta \ge \bar{\theta}\}$. One only needs to show that paths starting in $Q(\bar{\theta})$ must go to an arbitrarily small neighborhood of $\bar{\theta}$, and stay there. The details differ slightly from those in [6], and we comment briefly.

The proof in [6] of the last assertion depended on the fact that if there is a path of the ODE with value $\widehat{\theta}(t)$ at some time, where $\widehat{\theta}(t) \in Q(\bar{\theta}) - \{\bar{\theta}\}$, then $\dot{U}(\widehat{\theta}(t)) < 0$. Then $U(\cdot)$ can be used as a Liapunov function (together with the fact that all paths tend to the set $Q(\bar{\theta})$) to prove the final convergence and stability assertion. Thus, suppose that

$$\dot{U}(\widehat{\theta}) = \sum_i (\bar{h}_i(\widehat{\theta}) - \widehat{\theta}_i)/(d_i + \widehat{\theta}_i) \ge 0, \quad \text{for some } \widehat{\theta} \in Q(\bar{\theta}) - \{\bar{\theta}\} \tag{16}$$

Since $\widehat{\theta} \ge \bar{\theta}$ and $\widehat{\theta} \ne \bar{\theta}$, (16) implies that

$$\sum_i (\bar{h}_i(\widehat{\theta}) - \bar{\theta}_i)/(d_i + \bar{\theta}_i) > 0. \tag{17}$$

Now, consider the algorithm, started at $\bar{\theta}$, but with the power allocation rule at time n $\max_{u_i, i \le N} \left\{ \sum_i \log(1 + b_{i,n+1}u_i)/(d_i + \widehat{\theta}_i) \right\}$. I.e., the power allocation $\bar{u}(\widehat{\theta}, b_{n+1})$ is used. Modulo an error of order $O(\epsilon)t$, (12) and the maximizing property of $\bar{u}(\cdot)$ yields

$$\begin{aligned} &U(\theta^\epsilon(t)) - U(\bar{\theta}) \\ &= \epsilon \sum_i \sum_{l=0}^{[t/\epsilon]-1} \frac{\log(1 + b_{i,l+1}\bar{u}_i(\theta_l^\epsilon, b_l)) - \theta_{i,l}^\epsilon}{d_i + \theta_{i,l}^\epsilon} \ge \epsilon \sum_i \sum_{l=0}^{[t/\epsilon]-1} \frac{\log(1 + b_{i,l+1}\bar{u}_i(\widehat{\theta}, b_l)) - \theta_{i,l}^\epsilon}{d_i + \theta_{i,l}^\epsilon} \end{aligned} \tag{18}$$

where $[t/\epsilon]$ = integer part of t/ϵ, and θ_l^ϵ (with interpolation $\theta^\epsilon(\cdot)$) is from (12) under (11).

The SA arguments that led to Theorem 1, together with (18), imply that as $\epsilon \to 0$,

$$U(\theta(t)) - U(\bar{\theta}) = \int_0^t \sum_i \frac{\bar{h}_i(\theta(s)) - \theta_i(s)}{d_i + \theta_i(s)} ds \ge \int_0^t \sum_i \frac{\bar{h}_i(\widehat{\theta}) - \theta_i(s)}{d_i + \theta_i(s)} ds.$$

This, inequality (17), and the fact that $\bar{\theta}$ is an equilibrium point of the ODE, imply that

$$\dot{U}(\theta(t))|_{t=0} = \sum_i \frac{\bar{h}_i(\bar{\theta}) - \bar{\theta}_i}{d_i + \bar{\theta}_i} = 0 \ge \sum_i \frac{\bar{h}_i(\widehat{\theta}) - \bar{\theta}_i}{d_i + \bar{\theta}_i} > 0.$$

This contradiction to (16) implies that $\dot{U}(\theta) < 0$ for $\theta \in Q(\bar{\theta}) - \{\bar{\theta}\}$. Thus $\dot{U}(\theta(\cdot|\widetilde{\theta}))$ is strictly decreasing in $Q(\bar{\theta}) - \{\bar{\theta}\}$, which implies that a path starting at any $\theta(0) \in Q(\bar{\theta})$ must end up at $\bar{\theta}$, completing the proof that $\bar{\theta}$ is a unique asymptotically stable point.

3.2. Minimum Throughput Requirement. Now let there be constraints on the minimum throughput in that we desire $\theta_i \ge a_i$. For simplicity, work with the problem of Section 2. If the current θ_i value is less than a_i then user i needs to be given some advantage in that it should be selected even for some rates that are lower than what (5) would require. But how much advantage ? Overall efficiency is still a concern, and user i should not be selected if its current rate is too low. This reasoning argues against the use of hard constraints.s Consider the following "soft" approach. For each i, let $q_i(\theta_i)$ be a differentiable penalty function which is zero if $\theta_i > a_i$, negative for $\theta_i < a_i$, and whose derivative $q_{i,\theta_i}(\theta_i)$ is non decreasing in θ_i as θ_i decreases. For example, $q_i(\theta_i) = K_i \max\{0, |\theta_i - a_i|(\theta_i - a_i)\}, K_i > 0$. Analogously to the role of (9) in getting (5), let the allocation rule maximize the first order increment of $\sum[\log(\theta_i + d_i) + q_i(\theta_i)]$, a strictly concave function. [Or using whatever utility function replaces $log(\cdot)$.] This yields the rule

$$\arg\max \left\{ r_{i,n+1} \left[1/(\theta^\epsilon_{i,n} + d_i) + q_{i,\theta_i}(\theta^\epsilon_{i,n}) \right] \right\}, \tag{19}$$

The cooperative property holds, and Theorems 1, 3, and 4 hold under A1-A3.

The use of the penalty function allows flexibility near the boundary, since the choice will depend on both the current rates and by how much the constraint is violated. The use of penalty functions does not assure that the constraints will be satisfied. They might not be feasible. But it does give an advantage to users whose throughput is less than desired, and it puts a price on constraint violations.

3.3. Minimal SNR Required. Suppose that the additive noise at the receivers is white-Gaussian and that, due to the the decoding and error detection method, a user would not be selected unless a minimal SNR at the receiver of, say, κ, can be assured, and several users might be selected at the same time. If a user is selected, then the coding is based on the SNR being κ. Let $b_{i,n+1} = \alpha_{i,n+1}/N_{i,n+1}$ be as in Example 3.1. Then if user i is selected we must have $b_{i,n+1}u_i \geq \kappa$ (i.e., $u_i \geq \kappa/b_{i,n+1}$). With the minimal power applied to user i, the amount of data transmitted is a constant c_i. (5) will not yield the best allocation, since the power required to get the best ratio $c_i/[\theta_i + d_i]$ might yield a greater increment in the utility if it were split among several other users.

Let $I_{i,n+1}$ = indicator function of the event that user i is selected at n. Then

$$U(\theta^\epsilon_{n+1}) - U(\theta^\epsilon_n) = \epsilon \sum_i [c_i I_{i,n+1} - \theta^\epsilon_{i,n}]/[\theta^\epsilon_{i,n} + d_i] + O(\epsilon^2). \tag{20}$$

This tells us that the appropriate PFS algorithm is the maximization of

$$\sum_i c_i I_{i,n+1}/[\theta^\epsilon_{i,n} + d_i], \qquad \text{total power constraint } \kappa \sum_i I_{i,n+1}/b_{i,n+1} \leq \bar{u}. \tag{21}$$

The problem (21) is a form of the knapsack problem, since the $I_{i,n+1}$ are $\{0,1\}$-valued.[8] The cooperative property holds and, under A1–A3, Theorems 1, 3, and 4 hold.

A simpler form. A computationally simpler form drops the integer restriction and allows the $I_{i,n+1}$ be take arbitrary values in the interval $[0, 1]$. Then (21) is a linear program. Then solve (21) and select user i with the probability $I_{i,n+1}$. This might violate the instantaneous total power constraint $\bar{u}$, but it can be shown that the constraint will hold on the average, which is enough for many applications.[9]

[8] The throughput capabilities of PFS depends not only on the channels, but also on the processing capabilities of the scheduler. These are continually improving. In problems more complicated than that of Section 2, we should expect that the computation would be more complicated. While the exact solution of the general form of the knapsack problem can take too much time at present, in practice one would devise an algorithm yielding an acceptable approximation for the special problem at hand.

[9] Again, the LP algorithm would be specialized so that it efficiently yields an acceptable solution for the problem at hand.

3.4. Random Length Scheduling Intervals. There are recent applications (say, in EV-DO systems) where the number of slots required by a user selected at time n is random and not known at the time of selection. Consider the following example. The selected user transmits its data on the current slot. If acknowledgment is not received in the next slot, then it will use the next following slot to transmit additional data (perhaps repeating the original transmission, or sending additional parity bits). This can be repeated up to some finite maximum number of times. Then the number of slots required by the selected user is not known beforehand, and this must be taken into account in the selection rule. One application is where the channel for some users changes too fast to be well estimated, so an average SNR is used for them. But over any particular slot, the actual channel might be good or poor, and the scheme tries to adaptively account for this, using feedback from the receiver.

Under appropriate conditions an ODE can be obtained, but it will be hard to analyze. A simple approximation to this model, when one user is selected for each slot, supposes that the acknowledgments (if any) come at the beginning of the next slot. [I.e, the delay is ignored.] Then, if the nth selected user is i, it transmits for a bounded random number of slots $s^\epsilon_{i,n+1}$, which can depend on the current channel state j. Define $s^\epsilon_{n+1} = \sum_i I_{i,n+1} s^\epsilon_{i,n+1}$ and $\tau^\epsilon_n = \sum_{l=1}^n s^\epsilon_l$. Let $r_{i,\tau^\epsilon_{n+1}}$ denote the amount of data to be transmitted during the interval $[\tau^\epsilon_n, \tau^\epsilon_n + s^\epsilon_{n+1})$ if user i is selected. Since the time between selections is random, (3) needs to be modified by letting θ^ϵ_n denote the discounted throughput vector at the $(n+1)$st decision time, the end of the τ^ϵ_nth slot. The discounted form that converges to the average throughput per unit real time (as $\epsilon \to 0$) is

$$\theta^\epsilon_{i,n} = (1-\epsilon)^{\tau^\epsilon_n}\theta_{i,0} + \epsilon \sum_{l=1}^{n} (1-\epsilon)^{\tau^\epsilon_n - \tau^\epsilon_l} r_{i,\tau^\epsilon_l} I_{i,n}; \quad I_{i,n} = 1 \text{ if } i = n\text{th selection}, \tag{22}$$

$$\theta^\epsilon_{i,n+1} = \theta^\epsilon_{i,n} + \epsilon\left[r_{i,\tau^\epsilon_{n+1}} I_{i,n+1} - s^\epsilon_{n+1}\theta^\epsilon_{i,n}\right] + O(\epsilon^2). \tag{23}$$

Let j = channel state at τ^ϵ_n, when the $(n+1)$st scheduling decision is made. We have

$$U(\theta^\epsilon_{n+1}) - U(\theta^\epsilon_n) = \epsilon \sum_i \left[r_{i,\tau^\epsilon_{n+1}} I_{i,n+1} - s^\epsilon_{n+1}\theta^\epsilon_{i,n}\right] / \left[\theta^\epsilon_{i,n} + d_i\right] + O(\epsilon^2). \tag{24}$$

We cannot base the $(n+1)$st decision on maximizing (24), since s^ϵ_{n+1} is not known at that time. Suppose that the expectation of $s^\epsilon_{i,n+1}$, conditioned on both the data to the start of the $(n+1)$th scheduling interval and $r_{i,\tau^\epsilon_{n+1}}$ and j, depends only on the user choice i and on j at that time, and write it as $s_i(j)$. It is natural to replace $s^\epsilon_{i,n+1}$ in (24) by its conditional expectation $s_i(j)$. Then the appropriate PFS rule is (max over the $\{I_{i,n+1}\}$)

$$\max\left\{\sum_i \frac{r_{i,\tau^\epsilon_{n+1}} I_{i,n+1} - \left[\sum_l s_l(j) I_{l,n+1}\right]\theta^\epsilon_{i,n}}{\theta^\epsilon_{i,n} + d_i}\right\} = \max_i \left\{\frac{r_{i,\tau^\epsilon_{n+1}}}{\theta^\epsilon_{i,n} + d_i} - s_i(j)\sum_l \frac{\theta^\epsilon_{l.n}}{\theta^\epsilon_{l,n} + d_l}\right\}. \tag{25}$$

The choice of user i determines $s_i(j)$, the mean number of slots until the next selection, and the discounting from decision time τ^ϵ_n to decision time τ^ϵ_{n+1} in (22). For this reason, one cannot separate the users and evaluate them one at a time as in (5). Under conditions similar to A1–A2, Theorem 1 holds with the ODE being $\dot\theta_i = \bar h_i(\theta) - \bar s(\theta)\theta_i, \quad i \le N$, where $\bar s(\theta)$ denotes the mean number of slots required per scheduling period when the current state is θ, and $\bar h_i(\theta)$ is the average throughput for user i over a typical scheduling period. Since the throughput in (22) is updated at the iterate times, (26) is in the iterate time scale and not the real time scale. By a change of variables, the ODE in real time becomes $\dot{\theta}_i = \bar{h}_i(\theta)/\bar{s}(\theta) - \theta_i, \quad i \leq N$, where $\bar{h}_i(\theta)/\bar{s}(\theta)$ is the average throughput per unit time for user i when the state is θ. If the rate and interval length sequences are i.i.d., then the system can be shown to be cooperative and Theorems 3 and 4 hold. Results under weaker conditions are not yet available.

The method used in applications at this time uses (22), but the decisions use the rule

$$\arg\max \left\{ \sum_i \left[r_{i,\tau^\epsilon_{n+1}} I_{i,n+1} \right] / \left[\theta^\epsilon_{i,n} + d_i \right] \right\}. \tag{26}$$

This rule is simpler than (25). But it does not account for the expected duration of the required time. Theorem 1 holds under conditions similar to (A1)-(A2). Simulations for the i.i.d. case indicate that the limit point is unique. There are no associated optimality properties, since the rule does not maximize an increment of the utility function.

Consider the case where there are two users with i.i.d. rate and interval processes, whose rate processes are identically distributed, but the interval processes are not. Simulations imply that (26) gives each the same limit throughput. But, under (25), the user with the smaller mean interval has a larger throughput, as expected, and as is consistent with the concept of fair sharing, since its effective rate is higher. The average throughput seems higher under (25). So the new form (25) is preferable to (26). Little is known about the performance under either rule when the processes are correlated. Then (25) will in general not be optimal and there is room for improvement by taking a view that is longer than one interval at a time.

References

[1] R. Agrawal and V. Subramanian. Optimality of certain channel aware scheduling policies. Proc. 40th Allerton Conference on Communications, Control and Computing: 1532 - 1541, 2002.

[2] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushyana, and S. Viterbi. CDMA/HDR: a bandwidth efficient high speed wireless data service for nomadic users. *IEEE Comm. Magazine*, 38:70–77, 2000.

[3] S. Borst. User level aware performance of channel-aware scheduling algorithms in wireless data networks. In *Proc. Infocom, San Francisco*, New York, 2002. IEEE Press.

[4] J. M. Holtzman. Asymptotic analysis of proportional fair sharing. In *Proc. Personal, Indoor, and Mobile Radio Communications, Vol. 2*, pages 33–37, New York, 2001. IEEE Press.

[5] H.J. Kushner and P.A. Whiting. Asymptotic properties of proportional-fair sharing algorithms. In *Proc., 2002 Allerton Conf.*, Champaigne-Urbana, IL, 2002. Univ. of Illinois Press.

[6] H.J. Kushner and P.A. Whiting. Convergence of proportional-fair sharing algorithms under general conditions. *IEEE Trans. Wireless Communications*, 2003. To appear.

[7] H.J. Kushner and G. Yin. *Stochastic Approximation Algorithms and Applications.* Springer-Verlag, Berlin and New York, 1997. Second ed., 2003.

[8] J.G. Proakis. *Digital Communications, 4th Edition.* McGraw-Hill, New York, 2001.

[9] H.L. Smith. *Monotone Dynamical Systems: An Introduction to Competitive and Cooperative Systems, AMS Math. Surveys and Monographs, Vol. 41.* Amer. Math. Soc,., Providence RI, 1995.

[10] A.L. Stolyar. Multiuser throughput dynamics under proportional fair and other gradient based scheduling algorithms. Bell Labs. Report., 2002.

[11] P. Viswanath, D.N.C. Tse, and R. Laroia. Opportunistic beamforming using dumb antennas. *IEEE Trans. on Inf. Theory*, 48:1277–1294, 2002.