Convergence of Proportional-Fair Sharing Algorithms Under General Conditions

Harold J. Kushner*
Division of Applied Mathematics, Brown University, Providence RI 02912

Philip A. Whiting
Bell Labs, Lucent Technologies, 700 Mountain Ave,
Murray Hill NJ 07974-0636

Original version, July, 2002; Revision, Feb. 2003

Abstract

We are concerned with the allocation of the base station transmitter time in time varying mobile communications with many users who are transmitting data. Time is divided into small scheduling intervals, and the channel rates for the various users are available at the start of the intervals. Since the rates vary randomly, in selecting the current user there is a conflict between full use (by selecting the user with the highest current rate) and fairness (which entails consideration for users with poor throughput to date). The Proportional Fair Scheduler (PFS) of the Qualcomm High Data Rate (HDR) system and related algorithms are designed to deal with such conflicts. The aim here is to put such algorithms on a sure mathematical footing and analyze their behavior. The available analysis [6], while obtaining interesting information, does not address the actual convergence for arbitrarily many users under general conditions. Such algorithms are of the stochastic approximation type and results of stochastic approximation are used to analyze the long term properties. It is shown that the limiting behavior of the sample paths of the throughputs converges to the solution of an intuitively reasonable ordinary differential equation, which is akin to a mean

*Partially supported by Army Research Office Contract DAAD19-00-1-0549 and National Science Foundation Grant ECS ECS 0097447

flow. We show that the ODE has a unique equilibrium and that it is characterized as optimizing a concave utility function, which shows that PFS is not ad-hoc, but actually corresponds to a reasonable maximization problem. These results may be used to analyze the performance of PFS. The results depend on the fact that the mean ODE has a special form that arises in problems with certain types of competitive behavior. There is a large set of such algorithms, each one corresponding to a concave utility function. This set allows a choice of tradeoffs between the current rate and throughout. Extensions to multiple antenna and frequency systems are given. Finally, the infinite backlog assumption is dropped and the data is allowed to arrive at random. This complicates the analysis, but the same results hold.

1 Introduction: The Basic Algorithm

Consider the problem where there are a fixed number (N) of users competing to transmit data from a single base station to N mobile destinations, each moving independently of the others, and the possible rates of transmission of the individual users are randomly time varying. Time is divided into small scheduling intervals (called slots). Until further notice, in each interval one of the N users is chosen to transmit to its destination. If user i is selected in interval n, then it transmits $r_{i,n}$ units of data, where $\{r_{i,n}, n < \infty\}$ is a bounded (and usually correlated in n) random sequence, which might also be correlated among the i. They need only satisfy some mixing-type condition, specified in Section 2.

Motivation for the work comes from recent cellular systems, such as the Qualcomm High Data Rate system (the ISO 856 standard), which provide data connection via a common shared downlink onto which user transmissions are scheduled. In that system, access to the link is given one user at a time for a time slot of fixed duration of 1.67 ms. The decision as to which user is to be chosen for that slot is made on the basis of how much data (i.e, the rate) each could transmit over the interval, as well as the past history. These rates take into account the estimated SNR, which is determined via measurements based on a pilot signal. Since the time between measurement and prediction is short, fairly accurate rate predictions can be made. Scheduling decisions can take into account Rayleigh fading with a frequency of a few tens of Hertz.

The selection of the user at any time is based on a balance between the current possible rates and "fairness." One cannot choose the user with the highest rate at each slot, since users with lower SNRs will be starved. The question is how to "fair share." Fair sharing will lower the total throughput over the maximum possible, but it will provide more acceptable levels to users with poorer SNRs. The algorithm proposed by Qualcomm performs this sharing by comparing the given rate for each user with its average throughput to date, and selecting the one with the maximum ratio. This algorithm is known as *proportional* *fair sharing* (PFS) and is related to the fairness criterion given by Kelly [7] in allocating connections over multiple links on the Internet.

Until further notice, it is assumed that each user has an infinite backlog of data, the standard assumption in the literature to date. See [3, 11], where the channel behavior is stationary and ergodic. Such results indicate the gains that can be made by exploiting channel fluctuations that result from Rayleigh fading etc. Our assumptions will be weaker. Let the end of time slot (i.e., scheduling interval) n be called time n. At time n, the possible rates $\{r_{i,n+1}, i \leq N\}$ for the next time slot are known. Let $I_{i,n+1}$ be the indicator function of the event that user i is chosen at time n to transmit in slot $n+1$. One definition of the throughput for user i up to time n is the sample average

$$\theta_{i,n} = \sum_{l=1}^{n} r_{i,l} I_{i,l}/n, \tag{1.1}$$

where $I_{i,l+1} = 1$ if user i is chosen at time l and is zero otherwise. With the definition $\epsilon_n = 1/(n+1)$, (1.1) can be written in the recursive form (which defines Y_n)

$$\theta_{i,n+1} = \theta_{i,n} + \epsilon_n \left[I_{i,n+1} r_{i,n+1} - \theta_{i,n}\right] = \theta_{i,n} + \epsilon_n Y_{i,n}. \tag{1.2}$$

An alternative definition of throughput discounts past values of the $r_{i,n}$. For small positive ϵ, and the discount factor $1-\epsilon$, the discounted throughput is defined by

$$\theta^{\epsilon}_{i,n} = (1-\epsilon)^n \theta^{\epsilon}_{i,0} + \epsilon \sum_{l=1}^{n} (1-\epsilon)^{n-l} r_{i,l} I^{\epsilon}_{i,l}. \tag{1.3}$$

This can be written in the recursive form (which defines Y^{ϵ}_n)

$$\theta^{\epsilon}_{i,n+1} = \theta^{\epsilon}_{i,n} + \epsilon \left[I^{\epsilon}_{i,n+1} r_{i,n+1} - \theta^{\epsilon}_{i,n}\right] = \theta^{\epsilon}_{i,n} + \epsilon Y^{\epsilon}_{i,n}, \tag{1.4}$$

where $I^{\epsilon}_{i,n+1}$ is the indicator function of the event that user i is chosen at time n. The representations (1.2)-(1.4) allow arbitrary initial conditions, which might reflect some past history or bias. The recursive representation (1.2) allows the use of other values of the $\{\epsilon_n\}$. For example, they might go to zero more slowly than $1/n$, giving a weighting between those of (1.1) and (1.3). Owing to the boundedness of the $r_{i,n}$, the solutions to (1.2) and (1.4) are bounded.

The value of ϵ is chosen to balance the needs of estimating throughput (requiring a small value of ϵ) with the ability to track changes in the channel characteristics (requiring a larger value of ϵ). In general, ϵ should be chosen small enough so that it provides an acceptable measure of the throughput. A useful guideline is $\epsilon \times$[number of time slots by real time t] $= O(t)$.

The representations (1.2) and (1.4) are of the stochastic approximation form (see the comprehensive reference [9]), and the results of stochastic approximation theory will be used for their analysis. The possibility of using more general queues and the arbitrariness of the channel rate processes and of the ϵ_n illustrate the power of the approach. As will be seen, there are also convenient adaptations to multiple antenna/frequency systems. In addition, the assignment algorithm and convergence results will be seen to be typical of a large family of algorithms, each corresponding to some concave utility function, which in turn is actually maximized by the associated algorithm. Although the methods and conclusions exploit current results in stochastic approximation, they are far from obvious.

The original proportional-fair sharing algorithm chose the user that maximizes in[1]

$$\arg\max_{i \le N} \{ r_{i,n+1}/\theta_{i,n} \}, \tag{1.5}$$

or with $\theta^\epsilon_{i,n}$ replacing θ_n if (1.4) is used. When all of the current components $\theta_{i,n}, i \le N$, are very small, there is little sense in (1.5), since the current throughputs are all essentially zero and there is no reason to distinguish between them. We modify the algorithm slightly as follows. Let $d_i, i \le N$, be positive numbers, which can be as small as we wish. The chosen user at time n is that which maximizes in

$$\arg\max_{i \le N} \{ r_{i,n+1}/(d_i + \theta_{i,n}) \}, \tag{1.6}$$

or with θ^ϵ_n used if the algorithm is (1.4). In the event of ties, let us randomize among the possibilities. We choose randomization for resolving conflicts. But the end results are completely independent of how the conflicts are resolved.

Discretized or quantized rates. Suppose that the variables $r_{i,n}$ are the theoretical rates in that there is in principle a transmission scheme that could realize them, perhaps by adjusting the symbol interval and coding in each scheduling interval. In applications, it might be possible to transmit at only one of a discrete set of rates. The algorithms and results are readily adjusted to accommodate this need. Continue to make the assignments using (1.6), based on the values of the $r_{i,n+1}, i \le N$. But use the true transmitted rate, called $r^d_{n,i}$ in computing the throughput. Then (1.4) is replaced by

$$\theta^\epsilon_{i,n+1} = \theta^\epsilon_{i,n} + \epsilon \left(I^\epsilon_{i,n+1} r^d_{i,n+1} - \theta^\epsilon_{i,n} \right) = \theta^\epsilon_{i,n} + \epsilon Y^\epsilon_{i,n}, \tag{1.7}$$

Much is known about the behavior of algorithms such as (1.2) and (1.4); see, for example, [9, Chapters 8]. Under broad conditions, if the tracking parameter ϵ in (1.4) is small and

[1]The original algorithm, as applied to mobile communications, is due to D.N.C. Tse.

constant, or if n is large in (1.2), then the path converges to the solution to a deterministic ordinary differential equation, which is computed from the "mean" dynamics of the throughput process.

As we have just remarked, the path $\theta^\epsilon(\cdot)$ will essentially "follow" the solution to the ordinary differential equation (ODE). The ODE has a unique equilibrium point $\bar{\theta}$. The sample throughputs of (1.2) will converge to $\bar{\theta}$ for large n. With (1.4), the sample throughputs will congregate very close to $\bar{\theta}$ if ϵ is small. The existence of a unique equilibrium is of significance for the performance analysis of the proportional fair algorithm. This is because the equilibrium state will determine the throughput of each user and hence delay.

An interpretation of (1.6); maximizing a utility function. Define the utility function

$$U(\theta) = \sum_i \log(d_i + \theta_i). \tag{1.8}$$

An alternative view of (1.6) is that it maximizes $U(\theta_{n+1}) - U(\theta_n)$ to first order in the ϵ_n, as seen by the following argument. By a first order Taylor expansion,

$$U(\theta_{n+1}) - U(\theta_n) = \epsilon_n \sum_i \frac{r_{i,n+1} I_{i,n+1} - \theta_{i,n}}{d_i + \theta_{i,n}} + O(\epsilon_n^2). \tag{1.9}$$

Since $\sum_i I_{i,n+1} \equiv 1$, to maximize the first order term we must choose $I_{i,n+1}$ by (1.6). It will be shown in Section 3 that the rule (1.6) maximizes $\lim_n U(\theta_n)$ over all other admissible rules. See also the comments below (3.5) concerning the form (1.8). The function was not chosen apriori. It is the unique one associated with the original form of PFS.

The outline of the paper is as follows. Section 2 begins with some definitions. Then the assumptions are stated. They are stated in a general form, not only so that we can appeal to general results in stochastic approximation to facilitate the development, but also so that the basic structure that is required will be clear. The overall framework is quite flexible, and there are many useful variations. It is seen that the assumptions are quite reasonable. Then a weak convergence theorem for the iterative algorithms (1.2) and (1.4) is presented.

Theorems 2.2 and 2.3, which show that proportional fair has a unique stable point and optimizes the limiting utility $\lim U(\theta_n)$, are new. These results depend heavily on the structure of proportional fair and its relationship to what are known as monotone dynamical systems, as seen in Section 3. The existence of a unique equilibrium is of considerable significance for the performance analysis of PFS; otherwise, the usefulness would be questionable. To date, the primary alternative to such analysis has been simulation. Except for short transfers, the equilibrium state determines user throughput, hence delay. In Section 4, we give an example that illustrates the quality of the approach for modeling and analysis for the special case of Rayleigh fading.

Section 5 is devoted to extensions of the earlier results. The log concave utility is not essential. See Subsection 5.1. The results also hold if there are multiple channels with more than one user assigned at once. This allows applications such as scheduling from multiple transmit antennas, see Subsection 5.2. In Section 6, we drop the assumption of infinitely backlogged users. Suppose that (over the time period of interest), there are still a fixed number of users in the system, but they create data at random rates. This data is queued until transmitted. Under natural conditions, there is still a mean ODE which characterizes the flow of the algorithm. The queues themselves need to be brought into the analysis, since some might be empty at a scheduling time. This complicates the analysis, but the basic properties of and use of the mean ODE have not changed.

The convergence theorems which characterize the asymptotic behavior of the algorithms are given in Section 3, where it is also seen that argmax rule (1.6) also maximizes the utility function (1.8) for small ϵ and large n. The special case of constant rates with discontinuous dynamics is also discussed briefly. The utility function (1.8) plays no special role in the analysis. In Subsection 5.1, we see that other strictly concave utility functions can be used as well. This allows a choice of tradeoffs between the current rate and throughput in making the assignments. One advantage of the flexibility of the approach is the ease with which more complex systems can be treated. To illustrate this, in Subsection 5.2. we discuss the extension when there are multiple antennas or channels, and they can be assigned either individually or in some coordinated way.

Suppose that the users have a large, but not infinite, amount of data, and that new users can arrive from time to time. Then, during periods in which the number is fixed, the throughputs will follow the path of the mean ODE. For PFS to work well, one would need to choose an appropriate initial condition for any new user. It would be best if this were an estimate of the equilibrium value for that user in the current situation. For example, if we start a user with a throughput value of zero, then it will take slots even with very poor current rates. In Section 6, we drop the assumption of infinitely backlogged users. Suppose that (over the time period of interest), there are still a fixed number of users in the system, but they create data at random rates. This data is queued until transmitted. Under natural conditions, there is still a mean ODE which characterizes the flow of the algorithm. The queues themselves need to be brought into the analysis, since some might be empty at a scheduling time. This complicates the analysis, but the basic properties of and use of the mean ODE have not changed. Various extensions are reported in [8].

2 Assumptions and Convergence Theorems

First, some definitions will be given. Then we make a few comments concerning what is called weak convergence. The assumptions are then stated and discussed, after which the theorems are stated.

Definitions. Define the vectors $\theta_n = \{\theta_{i,n}, i \le N\}$, $\theta_n^\epsilon = \{\theta_{i,n}^\epsilon, i \le N\}$, and $R_n = \{r_{i,n}, i \le N\}$. The usual stochastic approximation asymptotic (or large time) analysis of the algorithms (1.2) and (1.4) uses continuous time interpolations. For each n, define the *shifted* process $\theta^n(\cdot)$ (with components $\theta_i^n(\cdot), i \le N$) by $\theta^n(0) = \theta_n$ and, for $l \ge 0$,

$$\theta^n(t) = \theta_{n+l} \text{ for } t \in \left[\sum_{k=n}^{n+l-1} \epsilon_k, \sum_{k=n}^{n+l} \epsilon_k\right),$$

where the empty sum is defined to be zero. Since the interpolated process $\theta^n(\cdot)$ starts at iterate n, the behavior of $\theta^n(\cdot)$ as $n \to \infty$ is that of θ_n as $n \to \infty$. Analogously, define the interpolated process $\theta^\epsilon(\cdot)$ (with components $\theta_i^\epsilon(\cdot), i \le N$) by $\theta^\epsilon(t) = \theta_n^\epsilon$ for $t \in [n\epsilon, n\epsilon + \epsilon)$.

Weak convergence. We work in a so-called weak convergence setup, which is the most powerful approach for the analysis of such algorithms [9, Chapter 8]. It is concerned with the characterization of the limits of the processes $\theta^n(\cdot)$ for large n and $\theta^\epsilon(\cdot)$ for large time. The details of the theory are not necessary for the development of the results. If we say that $\theta^\epsilon(\cdot)$ converges weakly to a process with constant value $\bar{\theta}$, it means that for large n, the paths of $\theta^\epsilon(\cdot)$ are very close to the point $\bar{\theta}$, with a high probability. The algorithm (1.4) requires the weak convergence view, since there will not be probability one convergence. One would not usually want to use (1.2) with the step size $\epsilon_n = 1/(n+1)$ since a few bad values of the noise in the early stages can mess up the behavior of the sample path for a long time to come, and such robustness considerations require that we have a larger discounting of past values than $\epsilon_n = 1/(n+1)$ provides. The weak convergence approach gives us much more flexibility than a probability one method, where that is possible. The discounted algorithm (1.4), or any form which is set up to allow tracking in the presence of time varying parameters, cannot converge w.p.1, so that a weak convergence analysis must always be done. More detail and discussion is in the comprehensive reference [9, Chapter 8].

Assumptions. The assumptions are quite weak. For convenience in applications, they are stated in a way that allows two options. Both options use (A2.0). The first option, which is (A2.1), is simpler to state and requires stationarity of $\{R_n\}$. The second, which is (A2.2), is implied by (A2.1) and allows more general behavior, such as deterministic cycling, etc. The conditions and the problem formulation are designed to take maximum advantage of the results in [9, Chapter 8]. Owing to the boundedness of the R_n, the conditions of the convergence theorems [9, Sections 8.2, 8.4] are implied by the forms used below. Let E_n

denote the expectation conditioned on $\{R_l, l \le n; \theta_0\}$, the data needed to calculate θ_n. The symbol r_i is used as the canonical value of the components $r_{i,n}$, and θ_i and θ_j are used as the canonical values of $\theta_{i,n}$ and $\theta_{j,n}$, resp. Let $I\!R_+^N$ denote the set of points $x \in I\!R^N$ with all components nonnegative. The requirements of the first stochastic approximation result, Theorem 2.1, are minimal. But, once the mean ODE is characterized, we need to know more about the set of possible limit points. The fact that the the limit point is unique and globally asymptotically stable is proved by using some results from dynamical systems theory.

A2.0. *Let ξ_n denote the past: $\{R_l : l \le n\}$. For each i, n, ξ_n,*

$$h_{i,n}(\theta, \xi_n) = E_n r_{i,n+1} I_{\{r_{i,n+1}/(d_i+\theta_i) \ge r_{j,n+1}/(d_j+\theta_j), j \ne i\}}$$

is continuous in $\theta \in I\!R_+^N$. Here θ is a parameter. Let $\delta > 0$ be arbitrary. Then in the set $\{\theta : \theta_i \ge \delta, i \le N\}$, the continuity is uniform in n and in ξ_n. The $r_{i,n}$ are uniformly bounded.

A2.1a. *$\{R_n, n < \infty\}$ is stationary. Define $\bar{h}_i(\cdot)$ by the stationary expectation:*

$$\bar{h}_i(\theta) = E r_i I_{\{r_i/(d_i+\theta_i) \ge r_j/(d_j+\theta_j), j \ne i\}}, \quad i \le N. \tag{2.1}$$

In (2.1), θ is considered fixed. Also,

$$\lim_{m,n \to \infty} \frac{1}{m} \sum_{l=n}^{n+m-1} \left[E_n r_{i,l+1} I_{\{r_{i,l+1}/(d_i+\theta_i) \ge r_{j,l+1}/(d_j+\theta_j), j \ne i\}} - \bar{h}_i(\theta) \right] = 0 \tag{2.2}$$

in the sense of probability. There are small positive δ and δ_1 such that

$$P\{r_{i,n}/d_i \ge r_{j,n}/(d_j - \delta) + \delta_1, j \ne i\} > 0, \quad i \le N. \tag{2.3}$$

A2.1b. *R_n has a bounded density.*

Comments on assumptions (A2.0) and (A2.1). The last part of (A2.1a) is innocuous and is used to assure that when a component θ_i is very small there is a nonzero chance that user i will be chosen, no matter what the values of the other components of θ. This is hardly a restriction. It guarantees that the mean rate function $\bar{h}_i(\theta)$ defined in (2.1) is positive when θ_i is small, The density assumption (A2.1b) is satisfied under standard physical assumptions. Indeed all the assumptions hold under Rayleigh fading if the channels are independent. The density condition is used only to show that the limit point is unique. Condition (2.2) is a very weak form of the law of large numbers, due to the use of the conditional expectation E_n. If the conditional expectation of the transmitted rate at time l, given the data to time n, is close to its stationary expectation for large positive $l - n$, then it holds. If the channel rate process is ergodic, then the condition holds even without the conditional expectation. So the combination of the effects of the conditional expectation and the division by m gives a very weak condition indeed. Condition (A2.0) asks that slight changes in θ_n would change the conditional (on data to the present) expectation of the next accepted rates only slightly. It is [9, condition (A2.3), Chapter 8]. The condition can be weakened by the approach taken in Section 6, but is good enough for typical applications where (A2.1b) holds.

On the smoothness of $\bar{h}(\cdot)$. By (A2.1b), $\bar{h}(\cdot)$ is Lipschitz continuous. To see this, consider the two dimensional case and let $p(\cdot)$ denote the density of R_n. Let $\theta \in I\!R_+^N$ and write $w = (d_1 + \theta_1)/(d_2 + \theta_2)$. Then

$$\bar{h}_1(\theta) = \int r_1 I_{\{r_1/r_2 \geq w\}} p(r_1, r_2) dr_1 dr_2,$$

which is Lipschitz continuous with respect to w, since the area of the region where the indicator is not zero is a differentiable function of w. The derivative of $\bar{h}_i(\cdot)$ will be continuous if $p(\cdot)$ is bounded and continuous.

A weaker set of conditions. The set of assumptions (A2.2a,b) is weaker than (A2.1a,b) in that it covers it and allows the distributions of the R_n to vary with n. For example, they can cycle in a deterministic way. A weak ergodicity property analogous to (2.2) is still required. In (A2.1), the condition (2.3) was sufficient to assure that the solution of (2.7) would have all components positive after an arbitrarily small time. We need a similar property when using (A2.2a,b) in order to avoid the possibility that some component will be zero in the limit due to degeneration of the rates. Because of the possibility of nonstationarity, a reasonably general sufficient condition cannot be expressed as simply as it was in (2.3). There are many sets of conditions which will assure it. For example, consider the random variables R_n^M obtained by randomizing among $\{R_{n+l}, l \leq M\}$ for some integer M. If R_n^M satisfies (2.3) uniformly in n, then the desired positivity will be assured. In order to avoid lots of examples, we simply write a sufficient condition (2.6) for what is required. Analogously to the situation with the set (A2.1), (A2.2a) is used in the basic stochastic approximation result, Theorem 2.1, and (A2.2b) is used to show that there is a unique limit point. Also, as noted above, the use of the conditional expectation in (2.5) gives a very weak form of the law of large numbers.

A2.2a. *The limits (stationary expectation used)*

$$\bar{h}_i(\theta) = \lim_{m,n\to\infty} \frac{1}{m} \sum_{l=n}^{n+m-1} E h_{i,l}(\theta, \xi_l) \tag{2.4}$$

exist and, for each θ,

$$\lim_{m,n\to\infty} \frac{1}{m} \sum_{l=n}^{n+m-1} E_n \left[h_{i,l}(\theta, \xi_l) - \bar{h}_i(\theta) \right] = 0, \quad i \leq N, \tag{2.5}$$

in probability. There are small positive δ, δ_1, *and an integer* M *such that*

$$\inf_n P\left\{r_{i,n+l}/d_i \geq r_{j,n+l}/(d_j - \delta) + \delta_1, j \neq i, \text{ for some } l \leq M\right\} > 0, \quad i \leq N. \tag{2.6}$$

A2.2b. *For any integer* M, *consider the random variables* R_n^M *obtained by randomizing among the* $\{R_{n+l}, l \leq M\}$. *There is* M *such that the* $R_n^M, n < \infty$, *have densities* $p_n^M(\cdot)$ *which are uniformly (in* n*) bounded.*

Note on quantized rates. If the *transmitted* rates are quantized as in (1.7), then redefine

$$h_{i,n}(\theta, \xi_n) = E_n r_{i,n+1}^d I_{\{r_{i,n+1}/(d_i+\theta_i) \geq r_{j,n+1}/(d_j+\theta_j), j \neq i\}},$$

$$\bar{h}_i(\theta) = E r_i^d I_{\{r_i/(d_i+\theta_i) \geq r_j/(d_j+\theta_j), j \neq i\}},$$

and in (2.2) use

$$E_n r_{i,l+1}^d I_{\{r_{i,l+1}/(d_i+\theta_i) \geq r_{j,l+1}/(d_j+\theta_j), j \neq i\}}.$$

The limit process and mean ODE. The next theorem is a standard result in stochastic approximation. It basically says that the limit points of the algorithms (1.2) and (1.4) are contained in those of the ODE (2.7).

Theorem 2.1. (This is [9, Theorems 2.2 and 2.3, Section 8.2].) *Assume the algorithm* (1.2), (A2.0), *and either* (A2.1a) *or* (A2.2a). *Then for any initial condition,* $\theta^n(\cdot)$ *converges weakly to the set of limit points of the solution of the ODE*

$$\dot{\theta}_i = \bar{h}_i(\theta) - \theta_i, \quad i \leq N. \tag{2.7}$$

The same conclusion holds if the ϵ_n *in* (1.2) *is replaced by a sequence* ϵ_n *such that* $\epsilon_n \to 0, \sum_n \epsilon_n = \infty$, *and where* ϵ_n *doesn't vary too fast in that for some sequence* $\alpha_n \to \infty$

$$\lim_n \sup_{0 \leq l \leq \alpha_n} |\epsilon_{n+l}/\epsilon_n - 1| = 0. \tag{2.8}$$

For algorithm (1.4), *the same conclusion holds for the sequence* $\theta^\epsilon(\epsilon q_\epsilon + \cdot)$ *for any sequence of integers* q_ϵ. *In particular, if* $\theta_0^\epsilon \to \theta(0)$, *then* $\theta^\epsilon(\cdot)$ *converges to the unique solution to* (2.7) *with initial condition* $\theta(0)$. *The conclusions hold if the discretized rates are used.*

Comment on the proof. The conditions of [9, Theorems 2.2 and 2.3, Section 8.2] are easily verified. The boundedness of R_n implies the uniform integrability needed in (the following cited conditions (conditions (A1.1)–(A1.9) are those in the reference) (A1.1) and (A1.8). The quantity β_n in (A1.5) is zero. In our case, we can suppose that ξ_n is defined on a compact sequence space, so (A1.7) is trivially satisfied. The averaging assumption (A1.9) is assured by either (2.2) or (2.5). The next two theorems will be proved in Section 3.

Theorem 2.2. *Assume algorithm* (1.2), (A2.0), *and either* (A2.1a,b) *or* (A2.2a,b). *The limit point* ($\bar{\theta}$) *of* (2.7) *is unique, irrespective of the initial condition. So the processes* $\theta^n(\cdot)$ *and* $\theta^\epsilon(\epsilon q_\epsilon + \cdot)$ *converge to* $\bar{\theta}$ *as* $n \to \infty$ *(resp., as* $\epsilon \to 0$ *and* $\epsilon q_\epsilon \to \infty$*, or as* $\epsilon \to 0$ *and* $t \to \infty$*). The conclusions hold if the discretized rates are used.*

Optimizing the utility function (1.8). Up to now, we have concentrated on the algorithm (1.6). Theorem 2.2 shows that there is a unique asymptotically stable limit point $\bar{\theta}$ of the ODE and algorithm. We have not addressed the optimality of the algorithm. We can see intuitively from the argument below (1.8) that (1.6) is, in the limit, a "steepest ascent" algorithm for a strictly concave utility function. The problem is that the allowed directions of ascent at each value of θ depend on θ. Hence there is no a priori guarantee of any type of maximization. The following theorem is one way of quantifying this idea.

Theorem 2.3. *Assume the conditions of Theorem* 2.2. *There is no assignment policy which yields a limit throughput* $\widetilde{\theta} \neq \bar{\theta}$ *such that* $U(\widetilde{\theta}) \geq U(\bar{\theta})$.

3 The Limit Point: Proof of Theorem 2.2

The proof depends on a monotonicity property of the solution to (2.7), which is fundamental to the analysis of a large class of stochastic algorithms which model competitive or cooperative behavior. The importance of these concepts was introduced in [5] and further discussion is in [1, 2]. First, some general facts from the theory of dynamical systems will be stated.

Definitions. For vectors $X, Y \in I\!R^N$, we write $X \geq Y$ (resp., $X > Y$) if $X_i \geq Y_i$ for all i (resp., and, in addition $X \neq Y$). If $X_i > Y_i$ for all i, then we write $X \gg Y$. Let H, K be subsets of $I\!R^N$. Then write $H \geq K$ (resp., $H \gg K$) if $X \geq Y$ (resp., $X \gg Y$) for all $X \in H, Y \in K$. Consider a dynamical system in $I\!R^N$:

$$\dot{x} = f(x), \qquad f(\cdot) \text{ Lipschitz continuous.} \tag{3.1}$$

Where helpful for clarity, we might write $x(t|y)$ for the solution to (3.1) when the initial condition is y, and use analogous notation for (2.7).

We say that the solution of (3.1) has a continuous dependence on perturbations if $x(t|y + \delta y)$ converges to $x(t|y)$ on each compact t-set as $\delta y \to 0$. The function $f(\cdot)$ is said to satisfy the *Kamke* (or simply the K-condition) if for any x, y and i, satisfying $x \leq y$ and $x_i = y_i$, we have $f_i(x) \leq f_i(y)$. In our case, $f(\theta) = \bar{h}(\theta) - \theta$, and the condition holds. The K-condition implies the following monotonicity result. Its proof in [10, Proposition 1.1] assumes continuous differentiability of $f(\cdot)$. But, it used only the K-condition, uniqueness of solutions, and the continuous dependence of the path on perturbations, all of which hold if $f(\cdot)$ is only Lipschitz continuous.

Theorem 3.1. [10, Proposition 1.1] Let $f(\cdot)$ be Lipschitz continuous and assume the K-condition. If $x(0) \le y(0)$ (resp., $<, \ll$, then $x(t|x(0)) \le x(t|y(0))$) (resp., $<, \ll$).

Proof of Theorem 2.2. By Theorem 2.1, we need only prove the assertions concerning the limit points of the paths of the mean ODE (2.7). Because of the boundedness of $\bar{h}(\cdot)$, the path $\theta(t|\theta(0))$ will be bounded, uniformly in the initial condition in any compact set, and all paths tend to some compact set as $t \to \infty$. Also, without loss of generality (say, by shifting the time origin), we can suppose where needed that there is $\delta > 0$ such that $\theta_i(t|\theta(0)) \ge \delta$ for all t. Since the path is initially monotone increasing (in each coordinate) when started near the origin (since $\bar{h}(\theta) - \theta \gg 0$ for θ near the origin), it follows from the monotonicity property (Theorem 3.1), that it will be monotonic (nondecreasing) in each coordinate for all t, for any initial condition sufficiently close to the origin. Thus, there is a unique limit point for the path $\theta(t|\theta(0))$ for each $\theta(0)$ near the origin.

Let $\bar{\theta}$ and $\widetilde{\theta}$ be two such limit points for paths corresponding to initial conditions $\bar{\theta}(0)$ and $\widetilde{\theta}(0)$, resp., arbitrarily close to the origin. By the monotonicity, $\theta(t|\theta(0)) \le \bar{\theta}$ for all $\theta(0)$ sufficiently close to the origin and all t. The next step is to show that $\bar{\theta} = \widetilde{\theta}$. All components of all paths that start very close to the origin are initially monotonically increasing. Thus for small enough $\bar{\theta}(0)$, we must have $\theta(t_0|\widetilde{\theta}(0)) \gg \bar{\theta}(0)$ for some $t_0 > 0$. Then, by monotonicity, $\theta(t+t_0|\widetilde{\theta}(0)) = \theta(t|\theta(t_0|\widetilde{\theta}(0))) \ge \theta(t|\bar{\theta}(0))$ for all $t \ge 0$. Thus, $\widetilde{\theta} \ge \bar{\theta}$. An analogous argument yields that $\bar{\theta} \ge \widetilde{\theta}$. We can conclude that there is a unique limit point, say $\bar{\theta}$, for all paths starting sufficiently close to the origin. Any limit point must be an equilibrium for (2.7) in that $h(\bar{\theta}) = \bar{\theta}$.

Now, consider the path starting at an arbitrary initial condition $\widehat{\theta} \le \bar{\theta}$. After some small time $t_0 > 0$, all components of the path will be positive. Hence there is $\theta(0) \gg 0$, and arbitrarily close to the origin, such that $\theta(t_0|\widehat{\theta}) \ge \theta(0)$. Then, the monotonicity argument of the above paragraph yields that $\theta(t + t_0|\widehat{\theta}) \ge \theta(t|\theta(0))$ for all t. Hence any limit point of $\theta(t|\widehat{\theta})$ must be no smaller than $\bar{\theta}$, the limit point of $\theta(t|\theta(0))$. But, by the monotonicity again, $\theta(t|\widehat{\theta}) \le \theta(t|\bar{\theta}) = \bar{\theta}$ for all t (the equality holds since $\bar{\theta}$ is an equilibrium point). We can conclude that $\theta(t|\widehat{\theta}) \to \bar{\theta}$ as $t \to \infty$.

Define the set $Q(\theta) = \{x : x \ge \theta\}$. Now, consider an arbitrary initial condition $\theta(0)$. The monotonicity argument can be used again to show that all limit points of the path are in $Q(\bar{\theta})$. It only remains to show that any path starting in $Q(\bar{\theta})$ must ultimately go to $\bar{\theta}$ also. So far, we have used only the monotonicity property and not any other aspect of the original stochastic approximation process that led to (2.7). The rest of the details involve the properties of the argmax rule and essentially standard stochastic approximation arguments. Suppose that there is a point $\widetilde{\theta} \in Q(\bar{\theta})$, $\widetilde{\theta} \ne \bar{\theta}$, such that

$$\dot{U}(\widetilde{\theta}) = \sum_i (\bar{h}_i(\widetilde{\theta}) - \widetilde{\theta}_i)/(d_i + \widetilde{\theta}_i) \ge 0. \tag{3.2}$$

Since $\widetilde{\theta} \geq \bar{\theta}$ and $\widetilde{\theta} \neq \bar{\theta}$, (3.2) implies that

$$\sum_i (\bar{h}_i(\widetilde{\theta}) - \bar{\theta}_i)/(d_i + \bar{\theta}_i) > 0. \tag{3.3}$$

Consider the algorithm (1.4) started at $\bar{\theta}$, but with the slot allocation rule

$$\arg\max_{i \leq N} \left\{ r_{i,n+1}/(d_i + \widetilde{\theta}_i) \right\}$$

used at time n. Let $\widetilde{I}^{\epsilon}_{i,n+1}$ denote the indicator function of the event that user i is chosen at time n. Modulo a second order error of order $O(\epsilon)t$, the expansion (1.9) and the maximizing property of I^{ϵ}_{n+1} (see (1.6)) yield

$$U(\theta^{\epsilon}(t)) - U(\bar{\theta}) = \epsilon \sum_i \sum_{l=0}^{[t/\epsilon]-1} \frac{r_{i,l+1} I^{\epsilon}_{i,l+1} - \theta^{\epsilon}_{i,l}}{d_i + \theta^{\epsilon}_{i,l}} \geq \epsilon \sum_i \sum_{l=0}^{[t/\epsilon]-1} \frac{r_{i,l+1} \widetilde{I}^{\epsilon}_{i,l+1} - \theta^{\epsilon}_{i,l}}{d_i + \theta^{\epsilon}_{i,l}} \tag{3.4}$$

where $[t/\epsilon]$ denotes the integer part of t/ϵ, and θ^{ϵ}_l (with interpolation $\theta^{\epsilon}(\cdot)$) is the solution to (1.4) under (1.6).

The stochastic approximation arguments that led to Theorem 2.1, together with (3.4), imply that as $\epsilon \to 0$ the limit $\theta(\cdot)$ satisfies

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

This, together with the inequality in (3.3) implies that

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

But the first sum is zero since $\bar{h}(\bar{\theta}) = \bar{\theta}$. Thus, we have a contradiction to (3.2) and can conclude that $\dot{U}(\theta) < 0$ for all $\theta \in Q(\bar{\theta}) - \bar{\theta}$. This implies that $\dot{U}(\theta(\cdot|\widetilde{\theta}))$ is strictly decreasing when the path is in $Q(\bar{\theta}) - \bar{\theta}$, which implies that any path starting at some $\theta(0) \in Q(\bar{\theta})$ must end up at $\bar{\theta}$. Thus, $\bar{\theta}$ is the unique limit point of (2.7), irrespective of the initial condition. Hence it is asymptotically stable. ■

A two-user example. Consider two independent users with received signal power determined by stationary Rayleigh fading (the Jakes model) and with constant external noise. Suppose further that their rate declarations are proportional to the absolute SNR, with mean rates $1/\beta_i, i = 1, 2$ respectively. Then the function $\bar{h}$ in (2.7) can be explicitly evaluated and we get

$$\begin{aligned} \dot{\theta}_1 &= \frac{1}{\beta_1} - \frac{\beta_1(d_1+\theta_1)^2}{\left(\beta_1(d_1+\theta_1)+\beta_2(d_2+\theta_2)\right)^2} - \theta_1, \\ \dot{\theta}_2 &= \frac{1}{\beta_2} - \frac{\beta_2(d_2+\theta_2)^2}{\left(\beta_1(d_1+\theta_1)+\beta_2(d_2+\theta_2)\right)^2} - \theta_2. \end{aligned} \tag{3.5}$$

For $d_i = 0, i = 1, 2$, it is readily shown that $\theta_i = 3/4 \cdot 1/\beta_i$ is a limit point. More generally $\theta_i = G(N)/N \cdot 1/\beta_i$ where $G(N) = \sum_{j=1}^{N} 1/j$. Note that this represents a scheduling gain, since with simple TDMA we would have $\theta_i = 1/[N\beta_i]$. This fact was also noted in [6], under the special conditions used there. The fact that the limit throughput is proportional to the mean 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 (1.6), the last fact being a consequence of the logarithmic utility function. In [4], under the same conditions it is shown that all users get the same fraction of slots, asymptotically. The argument supposes that there is an equilibrium point, a fact proved here.

Proof of Theorem 2.3: Optimality properties of the rule (1.6). We work with (1.4) for notational convenience. Suppose that there is an initial condition and a slot assignment policy that depends only on the available data and which attains some limit point $\widetilde{\theta} \neq \bar{\theta}$ such that $U(\widetilde{\theta}) \geq U(\bar{\theta})$; i.e., there is an admissible assignment sequence $\{\widetilde{I}_{i,n}; i \leq N,\ n < \infty\}$ under which the weak sense limit is $\widetilde{\theta}$ starting with the some given initial condition.

Now, consider the assignment algorithm which starts at the point $\bar{\theta}$ at time n, but uses the "tilde " strategy. This yields

$$\theta_i^\epsilon(n\epsilon + t) = (1-\epsilon)^{[t/\epsilon]}\bar{\theta}_i + \epsilon \sum_{l=1}^{[t/\epsilon]} (1-\epsilon)^{[t/\epsilon]-l} r_{i,n+l} \widetilde{I}_{i,n+l}. \tag{3.6}$$

Under (A2.0) and either (A2.1a) or (A2.2a), as $\epsilon \to 0$ and $n \to \infty$, this converges weakly to the process with values $e^{-t}\bar{\theta}_i + [1-e^{-t}]\widetilde{\theta}_i$. This, in turn, implies that the limit (as $\epsilon \to 0$) path goes in a straight line from $\bar{\theta}$ to $\widetilde{\theta} \neq \bar{\theta}$. Since $U(\widetilde{\theta}) \geq U(\bar{\theta})$, the rate of increase of the utility $U(\cdot)$ along that line is strictly positive (this is where the *strict* concavity of $U(\cdot)$ is used). An argument that exploits the maximizing property of the rule (1.6) and the equilibrium of $\bar{\theta}$ for (2.7), such as that used in the proof of Theorem 2.2, shows that it is impossible to have such a path in a neighborhood of $\bar{\theta}$. ■

4 Numerical Results

The graphs are from simulations based on Rayleigh fading, and the relation between the current rates and signal to noise ratios is taken from Table 1, which comes from [3]. Our first results depict the advantages to be gained by taking advantage of the current values of the time varying rates. In Figure 1, one set of curves corresponds to the transient behavior for three mobiles using table 1 and mean SNRs, -12dB,-2dB,-8dB, respectively, using algorithm (1.4). There are two sets of curves: those with solid lines and (the higher ones) those with dotted lines. The solid lines depict the throughputs if the SNRs (and hence the rates) are assumed to be constant at the average values. We take $\epsilon = 0.0001$. The value of ϵ is determined by a balance between what is considered a reasonable measure of discounted throughput and the desire to track changing conditions. If there are 1000 slots per second, then (roughly speaking) $\epsilon = 0.0001$ corresponds to a measure over about 10 seconds. Initially slots are offered only to mobile 2, with the throughputs for the other two mobile exponentially decaying. Also there are two "switching times". At first, the slots are equally divided between mobiles 2 and 3 $(0, 1/2, 1/2)$, then the slots are divided as $(1/3, 1/3, 1/3)$. (This behavior is generic for constant rates.)

The second set of curves (dotted lines/filled symbol) are obtained for Rayleigh fading with fading rate 6 Hz and the same mean SNRs. The true current rates are used. The significant gains in the throughput for all mobiles are evident. Since the dependence on rates is roughly linear on absolute SNR, it is expected the slots will be approximately evenly divided in equilibrium as the users all have exponential SNR distributions.

Example. Consider a case with two users with received signal power determined by a stationary Rayleigh fading process and with constant and white external noise. Suppose further that their rate declarations are proportional to the SNR, with mean rates $1/\beta_1, 1/\beta_2$ respectively. Using algorithm (1.4), the ODE is (3.5). For two such users and initial throughput 250.0, Figure 2 shows a sample path for the values of θ based on the proportional fair sharing algorithm with $\epsilon = 0.0001$, as well as a numerical solution to the corresponding ODE. The sample path rates were given via a Rayleigh fading simulator with $1/\beta_1 = 572$ bits/slot and $1/\beta_2 = 128$ bits/slot. The fading rates were taken as 60 Hz. With smaller values of ϵ and/or lower fading rates the sample paths fluctuate somewhat more about the solution to the ODE. In equilibrium the throughputs are, $\frac{3}{2} \cdot \frac{1}{2} \cdot 572 = 429, \frac{3}{2} \cdot \frac{1}{2} \cdot 128 = 96$.

The time constant for convergence is $1/\epsilon$ intervals and the results confirm convergence in a period of this order. The results also show the theoretical equilibrium being approached.

5 Extensions

5.1 Other Utilities and Allocation Rules

As noted in Sections 3 and 4. the algorithm (1.6) is based on the utility function $U(\theta) = \sum_i \log(d_i + \theta_i)$. Other strictly concave utility functions can be used as well and there are several that seem advantageous. One class is described next. Consider the utility

$$U(\theta) = \sum_i c_i(\theta_i + d_i)^\gamma, \quad 0 < \gamma < 1,\ c_i > 0,\ d_i > 0. \tag{5.1}$$

Then

$$\dot{U}(\theta) = \gamma \sum_i \frac{c_i \dot{\theta}_i}{(\theta_i + d_i)^{1-\gamma}}. \tag{5.2}$$

Thus the chosen user is

$$\arg\max_{i \le N} \left\{ \frac{c_i r_{i,n+1}}{(\theta_{i,n} + d_i)^{1-\gamma}} \right\}. \tag{5.3}$$

The ODE is $\dot{\theta}_i = Er_i I_{\{c_i r_i / c_j r_j \ge [(\theta_i + d_i)/\theta_j + d_j)]^{1-\gamma}, j \ne i\}} - \theta_i$. The rule (5.3) is not as sensitive to large values of θ_i as is (1.6). The analogs of Theorems 2.1–2.3 hold. Thus, there is a wide choice of useful and convergent algorithms which allow a variety of tradeoffs between the current rates and throughputs in making the assignment.

5.2 Extension to Multiple Channels and/or Antennas

Up to this point, there was only a single resource (say, a transmitter) to be assigned. There are similar algorithms and results when there are multiple resources to be assigned. To illustrate some of the possibilities, consider the following form, where there are two transmitters to be assigned, with possibly different locations and frequencies, but at the same base station. The associated channels will usually have different characteristics. For simplicity in exposition, we suppose that each user has an infinite backlog of data to be sent, and base the assignment rule on the utility (1.8) and the discounted throughput as in (1.4). In general, any number of antennas can be used.

One can allow many alternatives in the way that the transmitters are assigned. The examples are only intended to be illustrative of the possibilities that can be handled by the approach. In all cases to be discussed, it is assumed that the receiver at the mobiles are equipped to handle the method. The simplest method of assignment is to assign each of the two via an analog of (1.6) just as the single transmitter was assigned in Section 2. Then both might be assigned to one user, or only one might be assigned to each. The assignment algorithm is just (1.6), applied to each transmitter separately. Equivalently, one assigns so as to maximize the first order term in $U(\theta^\epsilon_{n+1}) - U(\theta^\epsilon_n)$, where $U(\cdot)$ is defined by (1.8). The assumptions are just those of Theorems 2.1–2.3, applied to the channels from each transmitter separately. In particular, there is a unique globally asymptotically stable limit point $\bar{\theta}$, and the argmax assignment algorithm maximizes the utility. The basic properties that were used in the proofs still hold. In particular, the mean ODE still satisfies the Kamke condition, and the argmax rule maximizes the increment in the utility to first order. Let $r_{ij,n}$ denote the canonical rates for user i in the channel from transmitter j at scheduling interval n. For the two user case, the iteration is

$$\theta^\epsilon_{1,n+1} = (1-\epsilon)\,\theta^\epsilon_{1,n} + \epsilon r_{11,n+1} I_{\{r_{11,n+1}/r_{21,n+1} \geq (d_1+\theta^\epsilon_{1,n})/(d_2+\theta^\epsilon_{2,n})\}}$$

$$+\epsilon r_{12,n+1} I_{\{r_{12,n+1}/r_{22,n+1} \geq (d_1+\theta^\epsilon_{1,n})/(d_2+\theta^\epsilon_{2,n})\}},$$

with the analogous formula for the other user.

There was no coordination between the assignments of the two transmitters in the method just discussed. Next consider an alternative that allows for more efficient use of the resource. We still allow the above choices, where each transmitter is assigned independently. But now we allow, in addition, the possibility of the two antennas being used in a coordinated way for the same user, with (for example) space-time coding. This simply adds another possible rate to be considered when using the argmax rule when making the assignment. Space-time coding is selected simply because it is one way of using both channels for the same user. The choice is still made for each scheduling interval, and might differ from interval to interval. The full assignment algorithm increases the channel capacity over what space-time coding used by itself could achieve. Let $r^c_{i,n}$ denote the rate using space-time coding when both channels are assigned to user i and let $I^{\epsilon,c}_{i,n}$ denote the indicator of this event. It is usually the case that $r^c_{i,n} \geq r_{i1,n} + r_{i2,n}$, and we make this assumption. For simplicity in the notation, the discussion is restricted to the case of two users.

Let $I^\epsilon_{ij,n}$ denote the indicator function of the event that user i is assigned to channel j in interval n but space-time coding is not used. Clearly, we need $\sum_i [I^\epsilon_{ij,n} + I^{\epsilon,c}_{i,n}] = 1$ for $j = 1, 2$, Then, to first order in ϵ, $U(\theta^\epsilon_{n+1}) - U(\theta^\epsilon_n)$ equals ϵ times:

$$\frac{r_{11,n+1} I^\epsilon_{11,n+1}}{d_1 + \theta^\epsilon_{1,n}} + \frac{r_{12,n+1} I^\epsilon_{12,n+1}}{d_1 + \theta^\epsilon_{1,n}} + \frac{r_{21,n+1} I^\epsilon_{21,n+1}}{d_2 + \theta^\epsilon_{2,n}} + \frac{r_{22,n+1} I^\epsilon_{22,n+1}}{d_2 + \theta^\epsilon_{2,n}}$$
$$+ \frac{r^c_{1,n+1} I^{\epsilon,c}_{1,n+1}}{d_1 + \theta^\epsilon_{1,n}} + \frac{r^c_{2,n+1} I^{\epsilon,c}_{2,n+1}}{d_2 + \theta^\epsilon_{2,n}} - \frac{\theta^\epsilon_{1,n}}{d_1 + \theta^\epsilon_{1,n}} - \frac{\theta^\epsilon_{2,n}}{d_2 + \theta^\epsilon_{2,n}}.$$

This yields a slightly more complicated form of the arg max rule of (1.6). One looks for the maximum of

$$\frac{r_{11,n+1}}{\theta^\epsilon_{1,n} + d_1} + \frac{r_{22,n+1}}{\theta^\epsilon_{2,n} + d_2},\ \frac{r_{12,n+1}}{\theta^\epsilon_{1,n} + d_1} + \frac{r_{21,n+1}}{\theta^\epsilon_{2,n} + d_2},\ \frac{r^c_{1,n+1}}{\theta^\epsilon_{1,n} + d_1},\ \frac{r^c_{2,n+1}}{\theta^\epsilon_{2,n} + d_2}.$$

Define $R_n = (r_{11,n}, r_{12,n}, r_{21,n}, r_{22,n}, r^c_{1,n}, r^c_{2,n})$. Suppose that R_n is bounded, stationary, and has a bounded density. Then, under the natural analogs of the other parts of (A2.0) and (A2.1a), the analysis of the resulting algorithm is similar to what was done for (1.2) and (2.4), and the analogs of Theorems 2.1–2.3 hold. Thus, even for this more complicated case, there is a unique limit point and the algorithm is a utility maximizer.

6 A General data Arrival Model

The formulation in Section 1 supposed that each user always has an infinite amount of data to be sent. This is, in fact, a shortcoming of the literature to date. Consider an alternative model, where there are still N users, but data arrives at random and is queued, awaiting transmission, and an arg max discipline such as (1.6) is used. We will confine our attention to the throughput as measured by (1.3), although (1.1) could be used as well. Let $Q^\epsilon_{i,n}$ denote the content of the queue for user i at time n. Define $Q^\epsilon_n = \{Q^\epsilon_{i,n}, i \leq N\}$ and $W^\epsilon_{i,n} = \min\{r_{i,n}, Q^\epsilon_{i,n-1}\}$. The decisions for the $(n+1)$st interval are still made at time n, when R_{n+1} is known. We will suppose that each queue i has a finite buffer of size B_i, and inputs to a full buffer are rejected and disappear from the system.

If $Q^\epsilon_{i,n} \geq r_{i,n+1}$ and queue i is selected, then an amount $r_{i,n+1}$ is transmitted and the queue decreases by that amount. But, if $0 < Q^\epsilon_{i,n} < r_{i,n+1}$, then one needs to modify the algorithm to reflect the fact that if queue i is selected, then the entire slot won't be filled. There are many ways of dealing with this problem. We will choose the approach of assuming that the decision is made on the basis of the current rates and the current value of the throughput, as in (1.6). More particularly, the decision at time l is

$$\arg \max_{i \leq N: Q^\epsilon_{i,l} > 0} \left\{ r_{i,l+1}/(d_i + \theta^\epsilon_{i,l}) \right\}. \tag{6.1}$$

The throughput is updated as (which defines Y^ϵ_n)

$$\theta^\epsilon_{i,n+1} = \theta^\epsilon_{i,n} + \epsilon \left[r_{i,n+1} I^\epsilon_{i,n+1} - \theta^\epsilon_{i,n} \right] I_{\{Q_{i,n} > 0\}} = \theta^\epsilon_{i,n} + \epsilon Y^\epsilon_{i,n}. \tag{6.2}$$

The motivation for (6.2) is that the scheduler will know only the rates and whether or not a queue is empty, but not the content of the queue. The scheme can be adjusted in many ways, taking into account whatever information is available. For example, after a queue reaches zero, there might be a latency period before it is polled again.[2] The queues evolve as

$$Q^\epsilon_{i,n+1} = Q^\epsilon_{i,n} + \delta A_{i,n+1} - W^\epsilon_{i,n+1} I^\epsilon_{i,n+1} - U^\epsilon_{i,n+1}, \tag{6.3}$$

where $\delta A_{i,n+1}$ denotes the number of arrivals to queue i in time slot $n+1$, and $U^\epsilon_{i,n+1}$ denotes the amount of data that was rejected due to a full buffer.

The model for the evolution of the queue and the $\{R_n\}$ will now be specified. We have in mind that there are essentially continuous arrivals to the queues, but at randomly time varying rates. The pair (Q^ϵ_n, r_{n+1}), together with δA_{n+1} and θ_n, determines $Q_{i,n+1}$. The pair

[2] Note added in proof: An alternative model that uses $\min\{r_{i,l+1}, Q_{i,l}\}$ in lieu of $r_{i,l+1}$ in (6.1) gives better results in simulations.

(Q^ϵ_n, r_{n+1}), together with θ^ϵ_n, determines θ^ϵ_{n+1}. The process $\{R_n, n < \infty\}$ does not depend on the evolution of the throughputs θ_n in that

$$P\left\{R_{n+1} \in \cdot | R_l, \theta^\epsilon_l, l \leq n\right\} = P\left\{R_{n+1} \in \cdot | R_l, l \leq n\right\}.$$

The contents of the queues, on the other hand, helps determine the evolution of the throughputs. The Q^ϵ_n (together with the R_n) play the role of "noise" and the θ^ϵ_n are the "states" of the system and we have what is called "state-dependent" noise [9, Chapters 6, 8] (at least the Q-component is state-dependent). The assumptions are stated in a somewhat abstract way since we wish to cover as many cases as possible within a single framework.

The main issue in the convergence proof concerns the fact that the evolution of the noise is determined by that of the throughputs, unlike the situation in Section 2. This complicates the averaging and use of conditions such as (2.2). But the fact that θ^ϵ_n varies slowly for small ϵ will help. Redefine the "memory" random variables to be $\xi^\epsilon_n = \{R_l, Q^\epsilon_l; l = n, n-1, \ldots\}$. Let E^ϵ_n denote the expectation, conditioned on the new ξ^ϵ_n. For each n and θ, define the *fixed-θ* process $\{\xi^\epsilon_l(\theta), l \geq n\}$ as follows: Suppose that after time n, we use the fixed value θ in (6.1) instead of the true current throughput. For $l > n$, we define it to be $Q^\epsilon_l(\theta)$, with "initial" condition $Q^\epsilon_n(\theta) = Q^\epsilon_n$, since the change is for times $l > n$ only. Define the functions

$$h_{i,n}(\theta, \xi^\epsilon_n) = E^\epsilon_n r_{i,n+1} I_{\left\{r_{i,n+1}/(d_i+\theta_i) \geq r_{j,n+1}/(d_j+\theta_j),\, j\neq i,\, Q^\epsilon_{j,n}>0\right\}} I_{\{Q^\epsilon_{i,n}>0\}}. \tag{6.4}$$

We will need the following conditions. They are quite weak, and their reasonableness will be seen by the example below.

(A6.1) There are Lipschitz continuous functions $\bar{h}_i(\cdot), i \leq N$, such that as $m \to \infty$

$$\frac{1}{m} E\left|\sum_{l=n}^{n+m-1} E^\epsilon_n\left[h_{i,l}(\theta, \xi^\epsilon_l(\theta)) - \bar{h}_i(\theta)\right]\right| \to 0 \tag{6.5}$$

uniformly in n, ϵ, and in θ in any compact set. There is $h_0 > 0$ such that $\bar{h}_i(\cdot) \geq h_0$ for small θ_i.

(A6.2) For each integer m,

$$\frac{1}{m} E\left|\sum_{l=n}^{n+m-1} E^\epsilon_n h_{i,l}(\theta^\epsilon_n, \xi^\epsilon_l(\theta^\epsilon_n)) - E^\epsilon_n h_{i,l}(\theta^\epsilon_l, \xi^\epsilon_l)\right| \to 0, \tag{6.6}$$

uniformly in n, as $\epsilon \to 0$.

(A6.3) For each integer m and as $|\widehat{\theta} - \widetilde{\theta}| \to 0$,

$$\frac{1}{m} E\left|\sum_{l=n}^{n+m-1} |E^\epsilon_n h_{i,l}(\widehat{\theta}, \xi^\epsilon_i(\widehat{\theta})) - E^\epsilon_n h_{i,l}(\widetilde{\theta}, \xi^\epsilon_i(\widetilde{\theta}))\right| \to 0, \tag{6.7}$$

uniformly in n and ϵ, where $\widetilde{\theta}$ and $\widehat{\theta}$ are in any compact set.

Example and discussion of the conditions. The Jakes model of Rayleigh fading is again covered, under reasonable conditions on the arrival process (e.g., as used in the example below). Conditions (A6.2) and (A6.3) are basically conditions on the sensitivity of the "conditional expectation of the amount transmitted" for user i to very small changes in the throughput that is used to make the decisions. Let us examine (A6.2) with $n < l \le n+m$. The difference between the terms $E_n^\epsilon h_{i,l}(\theta_n^\epsilon, \xi_l^\epsilon(\theta_n^\epsilon))$ and $E_n^\epsilon h_{i,l}(\theta_l^\epsilon, \xi_l^\epsilon)$ is that the second is the conditional expectation of the true amount transmitted for user i at l, given the data to n, and the first term is the conditional expectation of the amount that would have been transmitted if the rule (6.1) were used with $\theta = \theta_n^\epsilon$. However, over the time $[n, n+m)$, the change in θ_l^ϵ is bounded by ϵm, which goes to zero as $\epsilon \to 0$. Thus, the value of the state at the times of the summands in (6.6) is arbitrarily close to θ_n^ϵ, uniformly in n. Condition (A6.3) is similar. It says that if we make the decisions on $[n, n+m)$ always using either $\widehat{\theta}$ or $\widetilde{\theta}$, that are very close to one another, then the conditional mean amounts transmitted on that interval are also very close.

Let us illustrate the above comment via a simple example. Let $\{R_n\}$ and $\{\delta A_n\}$ be mutually independent, with the members of the latter sequence being mutually independent and identically distributed. Suppose that the component sequences $\{r_{i,n}, n < \infty\}$ are mutually independent in i. Let there be $0 \le \alpha_i < 1$, such that $r_{i,n+1} = \alpha_i r_{i,n} + \delta r_{i,n}$, where for each i, the $\delta r_{i,n}$ are mutually independent, identically distributed, bounded, and have a bounded and continuous density. Then, due to the Markov property, we can use $\xi_n^\epsilon = (Q_n, R_n)$. In (6.4), we have $r_{i,n+1} = \alpha_i r_{i,n} + \delta r_{i,n}$, where the $\delta r_{i,n}$ are bounded, mutually independent, and have a bounded density. This independence and density properties imply that small changes in the value of θ used in (6.4) changes the value of (6.4) only slightly, provided that the set of empty queues does not change. Furthermore, since the R_n are bounded, over any m iterates the value of θ_n^ϵ changes by at most $\epsilon K m$ for some constant K. By using the above facts and working forward from iterate to iterate, it can be seen that the probability that the decision at any $l \in [n, n+m]$ will be different (implying the possibility that the set of empty queues might change) for the two assignment methods goes to zero uniformly in n as $\epsilon \to 0$. Thus (A6.2) holds. Similarly, the probability that any assignment in $[n, n+m]$ based on the argmax rule using $\widehat{\theta}$ will differ from that based on $\widetilde{\theta}$ also goes to zero as $|\widehat{\theta} - \widetilde{\theta}| \to 0$. Thus (A6.3) holds.

Continuing with the example, (A6.1) is a weak ergodic condition. Let us make the decisions using the value θ for all $n \ge 0$, and not the true throughput. This yields the fixed-θ Markov process which we call $\xi_n(\theta) = (R_n, Q_n(\theta)$, for $n \ge 0$, with some arbitrary initial condition ξ_0. The process $\xi_n(\theta)$ has a unique invariant measure. Then, owing to the fact that the transition probability of the Markov process does not depend on time, (A6.1) says nothing more than that the conditional expectation (given ξ_0) of the throughput on $[0, n]$ corresponding to the continual use of the value θ, converges to the average value as $n \to \infty$. The convergence is uniform in the initial data and in the value of θ in any compact set.

Although the Rayleigh fading process is not Markovian, it has "mixing and density" properties that are similar to our Markov example.

Theorem 6.1. *Assume* (A6.1)–(A6.3). *Then the conclusions of Theorems* 2.1–2.3 *hold.*

Comment on the proof. Conditions (A6.1), (A6.2), and (A6.3), are close to conditions (A4.19), (A4.20), and (A4.21), resp., of [9, Theorem 4.6, Chapter 8], and can replace them in the proof of the cited theorem. Theorem 4.6 in [9, Chapter 8] is an analog for the state dependent noise case of the result cited in Theorem 2.1, and assures the conclusions of Theorem 2.1. The conclusions of Theorem 2.2 hold since they depend only on certain properties of the mean ODE, which hold in the present case. Theorem 2.3 depends on the stochastic approximation arguments which led to Theorem 2.1, the uniqueness of the limit point, and the strict concavity of the utility function, all of which hold in the present case.

References

[1] M. Benaïm and M.W. Hirsch. Mixed equilibria and dynamical systems arising from fictitious play in perturbed games. *Games and Economic Behavior*, 29:36–72, 1999.

[2] M. Benaïm and M.W. Hirsch. Stochastic approximation algorithms with constant step size whose average is cooperative. *Ann. Appl. Prob.*, 9:216–241, 1999.

[3] 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.

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

[5] M.W. Hirsch. Systems of differential equations that are competitive and cooperative: Convergence almost everywhere. *SIAM J. Math. Anal.*, 16:423–439, 1985.

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

[7] F. P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.

[8] H.J. Kushner and P.A. Whiting. Asymptotic Properties of Proportional-Fair Sharing Algorithms: Extensions of the Algorithm Proc., 2003 Allerton Conf., Univ. of Illinois Press, Champaigne-Urbana.

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

[10] H.L. Smith. *Monotone Dynamical Systems: An Introduction to Competitive and Cooperative Systems, AMS Math. Surveys and Monographs, Vol. 41.* American Mathematical Society, Providence RI, 1995.

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

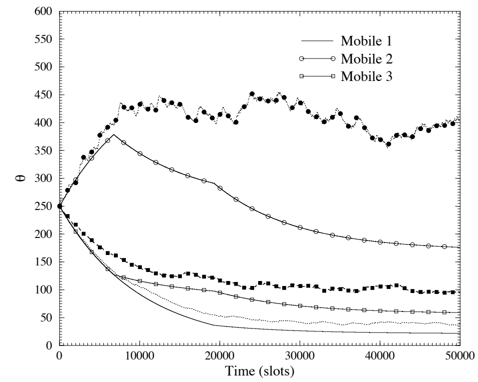

Figure 1: Time dependent behavior of Proportional Fair

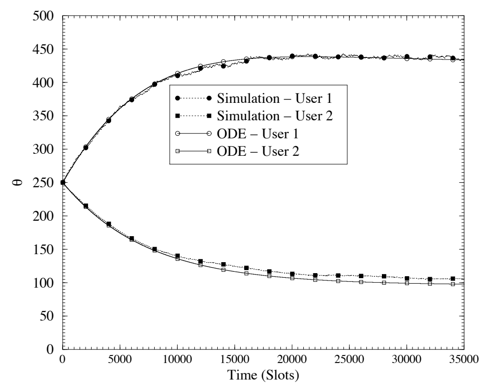

Figure 2: Sample path for θ and the Solution to the ODE

SNR	**-12.5**	**-9.5**	**-8.5**	**-6.5**	**-5.7**	**-4.0**
Rate	**0.0**	**38.4**	**76.8**	**102.6**	**153.6**	**204.8**
SNR	**-1.0**	**1.3**	**3.0**	**7.2**	**9.5**	**-**
Rate	**307.2**	**614.4**	**921.6**	**1228.8**	**1843.2**	**2457.6**

Table 1: Rate vs. SNR for 1% packet loss (taken from [3]