Adaptive Induced Fluctuations for Multiuser Diversity

Sujay Sanghavi and Bruce Hajek, *Fellow, IEEE*

***Abstract*— This paper investigates multiuser diversity in the context of cellular networks, with emphasis on the gains that can be achieved by adaptively inducing fluctuations in the environment.**

In a cellular system that at any time schedules for service the user with the best channel at that time, the expected service rate increases with the variability of the channel. Controlling the fluctuations using available feedback can further increase the expected service rate. This paper proposes a scheme for controlling fluctuations using only the feedback required to exploit multiuser diversity. Fluctuations are induced by introducing at the base station another transmit antenna that sends out the same signal but at a different phase from the first one, and then adaptively varying the phase difference.

The performance of the scheme (for the same total transmitted power) is evaluated when the users are infinitely back-logged or have finite queues, and when the channels are Rayleigh or Ricean distributed. Fairness issues and performance of the scheme under an additional fairness mechanism are also investigated in the context of users with finite queues. In all scenarios the performance is better when fluctuations are adaptively induced than when the fluctuations are randomly induced or not induced at all.

***Index Terms*— Multiuser diversity, scheduling for fading channels, multiple antennas, cellular networks.**

I. INTRODUCTION

IN some multiuser communication systems with many independently varying channels, service can be scheduled so that the user served at any particular time is the one seeing an exceptionally good channel at that time. This has been called *opportunistic scheduling* and the resulting gain has been referred to as the *multiuser diversity* gain. This gain increases as the number of users that can be possibly scheduled increases, since the system performance on average grows as $E[\max_i Capacity_i]$, the expected value of the maximum of the achievable data rates for the different users.

In such a scenario, for a given average value of the channel, a channel showing more variability would typically result in a higher multiuser diversity gain than one with less variability. For a wireless multiaccess scenario, Viswanath *et al.* [1] further suggest deliberately *inducing* variations in the channel by imposing random time-varying relative phases on the antennae of a transmit array. This increases the rate of channel variation, enabling the use of schemes harnessing multiuser diversity in environments where the time scale of variation of the channel is otherwise too slow for fairness over reasonable time scales.

The main purpose of this work is to investigate controlling the speed of the induced fluctuations, based on the available feedback (i.e., making the fluctuations *adaptive*). That is, during periods when the base station is sending at a relatively high data rate, the fluctuations are slowed down. Fluctuations are induced by changing the phase difference between the signals transmitted on different transmit antennae. Closed loop techniques to exploit transmit diversity are supported in some cellular wireless standards. For example, UMTS [7, pp. 114] supports two closed loop transmit diversity modes: one where the phase difference between the signals transmitted on the antennae can be controlled based on the feedback received from the mobiles, and another where the amplitude of the two signals can be controlled as well. The scheme proposed in this paper can be seen as a way of using this available feedback and control to leverage multiuser diversity in a better way than is possible with open loop schemes.

The idea of staying near the global maximum of a function using trajectories that decelerate where the function is large was investigated in a different, deterministic environment by Meerkov [2]. Besides enabling fairness for slow channels, an adaptive scheme also shows gains in average throughput for fast fading channels with memory, something that non-adaptive schemes cannot achieve.

Section II gives a brief overview and then a precise formulation of the system model and the problem. Section III details two upper bounds and one lower benchmark for judging the performance of any strategy for inducing fluctuations for the case when users are infinitely back-logged. Section IV describes the proposed scheme for adaptively inducing fluctuations. Sections V and VI present the performance of this proposed scheme for scenarios where users are all either infinitely backlogged or have queues, respectively. Section VI also investigates the effect on performance of additional mechanisms that ensure fairness.

II. SYSTEM MODEL AND PROBLEM FORMULATION

A. Brief Overview

Consider the basic cellular wireless downlink system shown in Figure 1(a), consisting of a central base station (BS) with one antenna communicating with many mobile users. Consider adding another antenna at the base station, as shown in Figure 1(b), which transmits the *same* signal as the first one but with a time varying relative phase that can be controlled. As a result of the addition, at any user the two signals constructively or destructively interfere, depending on the phase.

Manuscript received March 6, 2003; revised April 29, 2004; accepted June 26, 2005. The associate editor coordinating the review of this paper and approving it for publication was Q. Zhang. This work was supported by Grants NSF CCR 99-79381 and NSF ITR 00-85929.

The authors are with the University of Illinois at Urbana-Champlain (e-mail: sanghavi@uiuc.edu, hajek@comm.csl.uiuc.edu).

Digital Object Identifier 10.1109/TWC.2006.03115.

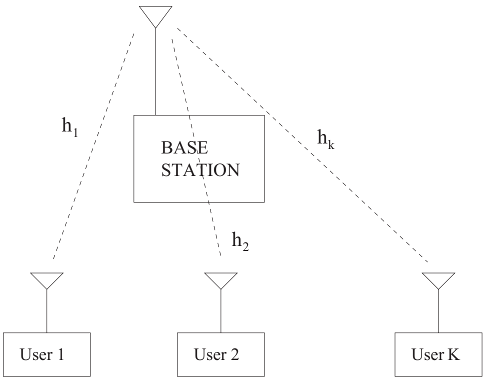

(a) Basic one antenna model.

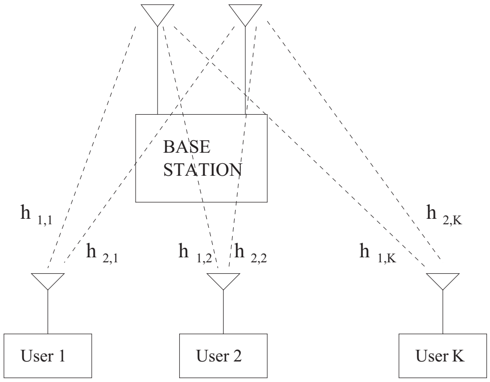

(b) Modified two antenna model.

Fig. 1. Two systems.

The interference patterns at the various users can be changed by changing the phase, thus effectively inducing controlled fluctuations in the channel. Further, if the users' channels possess memory, channel feedback can be utilized to control the phase so that the user scheduled for service experiences constructive interference and thus is served at a higher SNR (and thus a higher rate) as compared to the original single antenna system. This paper proposes and investigates a scheme that uses the channel feedback to control the phase to increase the throughput of the two antenna system. The rest of the section formalizes this system model.

B. Channel Models

The BS and users communicate over time varying fading AWGN channels. Time l is slotted, and the channel changes from slot to slot, while remaining constant within one slot. If there are N transmit antennae and K users, there are a total of NK channels in the system. We assume these channels are mutually independent.

The base station consists of N transmit antennae, all of which transmit the same $x(l)$ but with different phases $\{\theta_i(l)\}_{i=1}^{N}$. In the case of two antennae for instance the transmitted signals are $x(l)e^{j\theta_1(l)}$ and $x(l)e^{j\theta_2(l)}$. The phases $\{\theta_i(l)\}_1^M$ are in the control of the base station and can be varied according to some control strategy. Also, the total power transmitted at any particular time does not change, and is split among the antennae according to some (possibly time varying) splitting vector $\{\alpha_i(l)\}_{i=1}^{N}$ chosen so that $\sum_{i=1}^{N}\alpha_i(l)^2 = 1$.

Throughout this paper, "$\sim$" denotes equality of distribution. For $a \in \mathcal{C}$ and $b \in \mathcal{R}$, $b > 0$, let $\mathcal{CN}(a,b)$ denote the complex Gaussian distribution with mean a, with real and imaginary parts being independent Gaussians having variance $\frac{b}{2}$ each.

If $h_{n,k}(l)$ is the complex channel gain from the n^{th} BS antenna to the k^{th} user, the signal received at the k^{th} user at time l is

$$\begin{aligned} r_k(l) &= h_{1,k}(l)\alpha_1(l)e^{j\theta_1(l)}x(l) \\ &+ \ldots + h_{N,k}(l)\alpha_N(l)e^{j\theta_N(l)}x(l) + W_k(l) \quad (1) \end{aligned}$$

where $W_k(\cdot)$ is additive white Gaussian noise, assumed to be independent of the channel fading processes h with $W_k(\cdot) \sim \mathcal{CN}(0,1)$. It is also independent of the noise $W_j(\cdot)$ of any other user j.

The channel gains form a stationary proper complex random process, as defined in [3]. In particular $E[h_{i,k}^2(l)] = 0$ for all i, k and l. We consider two fading regimes - Rayleigh and Ricean - and Gauss Markov evolution in time[1]. For *Rayleigh channels* this means that for user k and BS antenna i,

$$h_{i,k}(l+1) = ah_{i,k}(l) + \sqrt{1-a^2}w_{i,k}(l) \qquad (2)$$

where $w_{i,k}(l)$ are mutually independent $\mathcal{CN}(0, 2\sigma_k^2)$ random variables and $a \leq 1$. This gives the process an autocorrelation of $R_{i,k}(l) = a^{|l|}(2\sigma_k^2)$, so a smaller a corresponds to a faster varying channel. If all the σ_k's have the same value then the users are said to have *symmetric channels*, else they are said to be *asymmetric*.

For *Ricean channels* we assume that the specular components $c_{i,k}$ remain invariant with time, and are all chosen to have the same deterministic magnitude and uniformly distributed arguments. The diffuse components $\tilde{h}_{i,k}$ evolve as in the case of Rayleigh fading, thus giving for all $i = 1,\ldots,N$ and $k = 1,\ldots,K$ the relations:

$$\begin{aligned} \tilde{h}_{i,k}(l+1) &= a\tilde{h}_{i,k}(l) + \sqrt{1-a^2}w_{i,k}(l) && (3) \\ h_{i,k}(l) &= \tilde{h}_{i,k}(l) + c_{i,k} && (4) \end{aligned}$$

Here the conditional autocovariance function (given the c's) is $C_{i,k}(l) = a^{|l|}(2\sigma^2)$, so a smaller a corresponds to a faster varying channel. The ratio of the powers in the specular and diffuse components is denoted by the Rice factor, defined as $\kappa \overset{\triangle}{=} \frac{\text{power in specular}}{\text{power in diffuse}}$. Here $\kappa = \frac{|c_{i,k}|^2}{2\sigma^2}$. The case $\kappa = 0$ thus corresponds to a Rayleigh channel, and $\kappa = \infty$ corresponds to a constant channel.

The Rayleigh channel gains have a circularly symmetric distribution, so $h'_{i,k} \overset{\triangle}{=} h_{i,k}e^{j\theta}$ has the same distribution as $h_{i,k}$ if θ does not depend on $h_{i,k}$. In the case of Ricean fading the diffuse components $\tilde{h}_{i,k}$ are circularly symmetric.

Simulations using the Jake's model for channel evolution over time are presented in the appendix. The channel model used and results are obtained are all described there.

C. Service Model

The service model consists of alternate test and service periods in each time slot, as shown in Figure 2. For each time slot l a particular choice of the transmit parameters is

---

[1] Simulations using the Jake's model are presented in the appendix

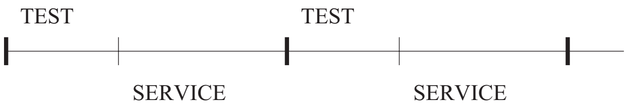

Fig. 2. The service structure.

fixed (say at $\{\theta_i(l), \alpha_i(l)\}$ for $i = 1, \ldots, N$). During the test period pilots are transmitted at a fixed known total transmit power according to the slot's transmit parameters, and each user records its received SNR, which it transmits back to the base station. The base station then selects the user reporting the best value of SNR and serves it exclusively for the length of the service period. Then, in the next time slot, new values of transmit parameters are chosen (say $\{\theta_i(l+1)\}_1^N$ and $\{\alpha_i(l+1)\}_1^N$) and a similar sequence of test and service is carried out for the new choice of parameters and user. The values of the transmit parameters for slot $l+1$ can depend on the values for slot l and on the SNR feedback from the various users that was received in the test period of slot l.

The throughput achievable in a particular time slot depends on the SNR of the user scheduled in that slot. In this paper we will consider the *Low SNR regime* where the throughput is assumed to be proportional to the SNR itself and the *High SNR regime* where it is proportional to $\log\left(1 + \frac{SNR}{SNR_{ref}}\right)$ where SNR_{ref} is the average SNR seen by one user if there was only one antenna transmitting all the power.

D. User Models

We will be working alternately with two types of users. The corresponding performance measures are detailed in the next section.

1) *File Transfer User:* A user that always has an infinite amount of data to send, and so the data it transmits in a particular time slot is limited only by the SNR seen by the user in that time slot. If a station is served during a particular time slot, it will be able to send an amount of data that is a function of the SNR it sees. The function depends on the SNR regime it is operating in, as described in Section II-C above.
2) *Real Time User:* A user that at any time has a finite amount of data to send, which it stores in its own queue. Data arrives in the queue according to some arrival process – in our simulations we assume that in each slot an arrival of m data units occurs at a user with probability p. The arrivals are independent across time and across users. The amount of data that can be sent by a real-time user when it is scheduled for service is the minimum of the amount of data that can be supported by its SNR and the amount of data it has available for transmission at the time. This minimum is the *potential amount of data* that user can transmit in that slot. In each slot the base station serves the user with the highest potential amount of data to transmit. Thus, for each user the feedback at each time consists of both the SNR as well as the amount of data in its queue.

E. Problem Formulation

Varying the transmit parameters $\{\theta, \alpha\}$ can be interpreted as inducing fluctuations in the channel seen by the users, and it is our objective to find an adaptive strategy that varies $\{\theta, \alpha\}$ so as to get better performance. We now state the problem formally for the restricted case when there are only two transmit antennae at the base station, and we have to find a good adaptive strategy for varying the relative phase θ between them. The transmitted power is split equally between the two antennae, so θ is the only parameter that is controlled. From (1), if $\theta(l)$ is the value of the parameter at time l, the SNR received by the k^{th} user for this case is

$$SNR_{k,l}(\theta(l)) \triangleq \frac{1}{2}\left|h_{1,k}(l) + e^{j\theta(l)}h_{2,k}(l)\right|^2 \qquad (5)$$

As mentioned, the performance measures for the two types of users means there are two separate problem statements for the two scenarios.

The file transfer scenario: All the users are of file transfer type and the system tries to maximize the total throughput. Since maximizing the service rate at each time is the only goal, the SNR at which service is given at time l is the best SNR over all users:

$$SNR_l^*(\theta(l)) \triangleq \max_k SNR_{k,l}(\theta(l)) \qquad (6)$$

The performance of any trajectory $\theta(\cdot)$ is the achievable throughput per slot, averaged over all possible channel realizations. For a scheme π, if the resulting (random, possibly channel-dependent) trajectory is $\theta_\pi(\cdot)$, the performance criterion for the scheme is the average achievable throughput per slot, denoted by f:

$$f(\pi) \triangleq E\left[\frac{1}{L}\sum_{l=1}^{L} SNR_l^*(\theta_\pi(l))\right] \qquad (7)$$

where L is the number of slots of interest, chosen so as to be much longer than the channel coherence time $\frac{1}{a}$. The expectation is taken over all channel realizations as well as over any possible randomness in the scheme itself. This includes the case where the trajectory $\theta_\pi(\cdot)$ may depend on the channel realization.

Similarly, for the high SNR regime it is given by

$$f(\pi) \triangleq E\left[\frac{1}{L}\sum_{l=1}^{L}\log\left(1 + \frac{SNR_l^*(\theta_\pi(l))}{SNR_{ref}}\right)\right] \qquad (8)$$

The aim of our work is to find a scheme π that only uses the SNR feedback described in Sections II-C and II-D to give an improvement in the value of f over a non-adaptive trajectory.

The real time scenario: All the users are of real time type. Here throughput is interpreted as the maximum arrival rate that can be supported while still keeping the queues of the users stable. Besides throughput the system also tries to minimize the average delay experienced for stable arrival rates. The problem is formulated below only for the low SNR regime.

For a given trajectory $\theta(\cdot)$, let $q_{k,l}(\theta)$ be the amount of data in the queue of the k^{th} user at time l. At time l the user chosen for service is the one with the highest potential amount of data to transmit: $\min\{SNR_{k,l}(\theta(l)), q_{k,l}(\theta)\}$. The average delay for this scenario can be found by Little's law. Let L be the length of time of interest, again chosen to be much greater than the coherence time $\frac{1}{a}$. For a scheme π, if the resulting (random, possibly channel-dependent) trajectory is $\theta_\pi(\cdot)$, the performance criterion for the scheme f is the average delay for all users, given by:

$$\text{Delay}(\pi) = E\left[\frac{\frac{1}{L}\sum_{l=1}^{L}\sum_{k=1}^{K} q_{k,l}(\theta_\pi)}{Kmp}\right] \tag{9}$$

where the expectation is taken over the realizations of the channel as well as any randomness the scheme might contain. The denominator Kmp is the average arrival rate of data into the *system*, i.e. all the users.

An arrival rate Kmp to the system is said to be unstable for a particular scheme if as $L \to \infty$ the mean delay converges to infinity. However, since we work with a large but finite horizon, we assume there is some value of the delay that is the threshold - an arrival rate with a resulting delay more than the threshold is said to be unstable for that scheme.

The aim of our work is to find a scheme for varying the trajectory of θ that only uses the SNR and queue feedback described in Section II-C and Section II-D to give an improvement in the delay as well as in the maximum stabilizable rate, over a non-adaptive trajectory.

III. BOUNDS FOR FILE TRANSFER USERS

In this section we consider a benchmark against which any scheme to induce fluctuations in the file transfer scenario (within the framework defined in Section II) can be compared. We also give upper bounds on the performance of any scheme. All users are assumed to have independent and identically distributed channels.

A. Benchmark - Open Loop Schemes

A scheme for selecting the trajectory is called "open loop" if it does not use channel feedback. Typical examples of this are keeping θ fixed, moving it with a constant velocity, or making it follow a random process that is independent of the channel. Open loop schemes serve as a benchmark for comparing the performance of strategies that control fluctuations.

Proposition:

(a) *If all the users have channels that display circular symmetry (e.g. Rayleigh channels), all open loop schemes give the same mean throughput for a given number of antennae.*

(b) *Further, in the case of pure Rayleigh fading, the mean throughput of all open loop schemes is the same, irrespective of the number of transmit antennae N.*

Proof:

Consider the case when there are N antennae transmitting at $\{\theta_i\}_1^N$ in a particular time slot to a user k over the channels $\{h_{i,k}\}_{i=1}^N$, each of which is circularly symmetric. The SNR $X_k^{(N)}$ seen by user k is

$$X_k^{(N)} = \left|\alpha_1 h_{1,k} e^{j\theta_1} + \alpha_2 e^{j\theta_2} h_{2,k} + \ldots + \alpha_N e^{j\theta_N} h_{N,k}\right|^2$$

Since $h_{i,k}$ is circularly symmetric and the value of θ_i does not depend on $h_{i,k}$, if $h'_{i,k} \overset{\triangle}{=} h_{i,k}e^{j\theta_i}$ then $h'_{i,k} \sim h_{i,k}$ and so

$$X_k^{(N)} \sim |\alpha_1 h_{1,k} + \alpha_2 h_{2,k} + \ldots + \alpha_N h_{N,k}|^2$$

Thus the distribution of $X_k^{(N)}$ does not depend on the θ's, i.e. the SNR seen by any user will have the same distribution for any open loop scheme. This implies that the expected (or long run) throughput will be the same for any open loop variation of the $\{\theta_i\}_1^N$. This proves part (a) of the proposition.

For the proof of (b), since $\sum_{i=1}^N \alpha_i^2 = 1$, the SNR $X_k^{(N)}$ has the same distribution as $|h_{i,k}|^2$; namely it is distributed as the magnitude of a $\mathcal{CN}(0, 2\sigma^2)$ random variable , which is exponential with mean $2\sigma^2$. In particular, this distribution does not depend on N. This proves part (b) of the proposition. □

Note that in the case when there is only one transmit antenna at the base station and Rayleigh fading, the performance of a system that at any time serves the best user is the same as the performance of the open loop θ variation. Thus the curves in the figures of sections V and VI labelled as "open loop θ" correspond to the case when there is only one antenna at the base station, *and* for the case when there are multiple antennas but open loop variations.

If the channels are Rayleigh fading channels, the SNR Z at which service occurs in the particular time slot for a system with K users is given by $Z \overset{\triangle}{=} \max_{1\leq k\leq K} X_k$, where the X_k's are independent exponentially distributed random variables with mean $2\sigma^2$. As shown in Appendix A the mean $E[Z]$ and standard deviation σ_Z are given by

$$\begin{aligned} E[Z] &= 2\sigma^2\left(1 + \frac{1}{2} + \ldots + \frac{1}{K}\right) \\ &\approx 2\sigma^2 \log K \text{ (for large } K) \\ \sigma_Z &= 2\sigma^2\sqrt{1 + \frac{1}{2^2} + \ldots + \frac{1}{K^2}} < \sigma^2\pi\sqrt{\frac{2}{3}} \end{aligned}$$

Note that the average throughput f, given by equations (7) or (8), of any open loop scheme will not depend on the speed a of the channel. The fact that inducing fluctuations in an open loop manner does not lead to an increase in throughput for independent fast-fading Rayleigh channels was noted in [1]. However, as also noted in [1], open loop induced fluctuations give significant gains in Rician and constant fading environments when fairness is taken in to account.

B. Upper Bounds - Beamforming to the Best User

We now present two upper bounds on the throughput of adaptive schemes, for the cases when the relative power split between antennae can and cannot be varied. The bounds are based on beamforming to the best user, which is possible only with perfect channel feedback.

Proposition: *For K users and N BS antennae having the propagation model described in (1),*

(a) *The SNR at which service can be given $Z_a^{(N,K)}$ is bounded by*

$$Z_a^{(N,K)} \leq \max_{1\leq k\leq K} |h_{1,k}|^2 + |h_{2,k}|^2 + \ldots + |h_{N,k}|^2 \tag{10}$$

with equality iff both the phases θ and the power splitting α can be optimally controlled using perfect knowledge of both the phase and the magnitude of each $h_{n,k}$.

(b) *Further, if the power is split equally among the transmit antennae, the SNR $Z_b^{(N,K)}$ at which service can be given is bounded by*

$$Z_b^{(N,K)} \leq \max_{1\leq k\leq K} \frac{1}{N}\left(|h_{1,k}| + |h_{2,k}| + \ldots + |h_{N,k}|\right)^2 \quad (11)$$

with equality iff the phases θ are optimally controlled using perfect knowledge of the phase of each $h_{n,k}$.

Equations (10) and (11) thus represent upper bounds for the performance of any scheme to induce fluctuations in their respective control scenarios, with or without perfect knowledge of the channel. These bounds are independent of the distribution of the channel gains h.

Proof:

For the propagation model of equation (1), the SNR $X_k^{(N)}$ seen by the k^{th} user in a particular time slot is given by

$$\begin{aligned} X_k^{(N)} &= \left|\alpha_1 h_{1,k} e^{j\theta_1} + \alpha_2 e^{j\theta_2} h_{2,k} + \ldots + \alpha_N e^{j\theta_N} h_{N,k}\right|^2 \\ &= \langle \underline{\rho}, \underline{\mathbf{h}_k}\rangle \end{aligned}$$

where

$$\begin{aligned} \underline{\rho} &\stackrel{\triangle}{=} \left[\alpha_1 e^{j\theta_1}\,,\, \alpha_2 e^{j\theta_2}\,,\, \ldots,\, \alpha_N e^{j\theta_N}\right] \\ \underline{\mathbf{h}_k} &\stackrel{\triangle}{=} \left[h_{1,k}\,,\, h_{2,k}, \ldots,\, h_{N,k}\right] \end{aligned}$$

By the Cauchy-Schwarz inequality, the choice of $\underline{\rho}$ that maximizes the $X_k^{(N)}$ is

$$\underline{\rho}^* = \frac{\underline{\mathbf{h}_k}^*}{\|\underline{\mathbf{h}_k}\|}$$

since there is a constraint on the total power, $\sum_1^N \alpha_i^2 = 1$ or equivalently $\|\underline{\rho}\|^2 = 1$. Under this choice, called beamforming, the SNR seen by user k is $|h_{1,k}|^2 + |h_{2,k}|^2 + \ldots + |h_{N,k}|^2$ and hence, in the case when the α_is can be controlled, the system will serve the user with the best SNR given by $Z_a^{(N,K)}$. This proves (10).

If the power splits α_i are not controllable but are fixed at $\frac{1}{\sqrt{N}}$ for each i, the optimal choices for the phase are still given by $\theta_i^* = -arg(h_{i,k})$, and with this choice the SNR seen by user k is $\frac{1}{N}\left(|h_{1,k}| + |h_{2,k}| + \ldots + |h_{N,k}|\right)^2$ and thus service will occur at the SNR given by $Z_b^{(N,K)}$. This proves (11). □

If the channels of the users are independent and identically distributed, the throughput when beamforming is used is N times higher than in the one-antenna case, which for Rayleigh fading is also the throughput under open-loop schemes for N antennae. There is thus a significant gain to be had from the proper control of the transmit parameters.

It is shown in Section I-C of the Appendix that the expected system throughput under beamforming in phase and power split, $E[Z_a^{*(2,K)}]$, for the case of 2 BS antennae can be approximated as follows:

$$\begin{aligned} &\max_{2\leq K\leq 1000} \left| E[Z_a^{*(2,K)}] - 2\sigma^2(\log(K+1) - \log(1+\log K) - 1)\right| \\ &\qquad \leq \quad 0.1248\sigma^2 \end{aligned}$$

A method for numerically calculating the exact value of $E[Z_a^{*(N,K)}]$ is also detailed Section I-C. The upper bounds presented here have been plotted in Section V.

As for the case of the throughput of open loop strategies, note that the upper bounds on throughput presented above do not depend on the speed a of the user channels.

IV. ADAPTIVE SPEED CONTROL

This section details the proposed scheme to induce fluctuations in the channel by adaptively varying the transmit parameters $\{\theta_i, \alpha_i\}(l)$ using the available feedback, assumed to be the user SNRs. The scheme presented here is for the restricted two antenna case where there is control only over the relative phase θ. The exact problem for this case has been formulated in Section II-E. However, the mechanisms given here may be generalized to handle more antennae or controlled power splitting.

For the file transfer scenario, the scheme we propose for controlling the relative phase $\theta(l)$ between the two transmit antennae is given by the update equation

$$\begin{aligned} \Delta\theta(l+1) &\stackrel{\triangle}{=} \theta(l+1) - \theta(l) \\ &= Be^{-\gamma SNR_l^*(\theta(l))} \end{aligned} \quad (12)$$

where the best value $SNR_l^*(\theta(l))$ is as defined by (5) and (6) and B and γ are the two parameters of the system. A larger SNR_l^* yields a smaller step size, so θ dwells longer in good regions of the phase space. At any time the user scheduled for service is the one with the highest SNR at that time.

The scheme for the real time scenario is similar, except that it now depends on the potential amount of data that can be transmitted and not just the SNR. The update rule is thus the same as in (12), but with

$$SNR_l^*(\theta(l)) = \max_k \left[\min\left\{SNR_{k,l}(\theta(l))\,,\, q_{k,l}(\theta)\right\}\right]$$

The user with the highest potential amount of data to transmit is the one that is served.

In both cases a suitable choice of B and γ depends on the average SNR, the type of channel (Rayleigh or Ricean), the number of users in the system, and how fast the channel is varying. Note that $\theta(l)$ always cycles around its phase space of $[0, 2\pi]$ in the same direction, but with varying speed. Also note that the rate of variation is decreasing in γ and increasing in B. The feedback required for the implementation of this scheme is thus the same as that required for opportunistic strategies with open loop trajectories for θ.

Some motivation for the form of the scheme is in order. For a given time slot, there is an optimum value of θ which corresponds to beamforming to the user with the best channel at that time. Any scheme for inducing fluctuations in the scenario under consideration has to have two properties: it should be able to search around the phase space $\theta \in [0, 2\pi]$ for the best value, and it should be able to spend a substantial

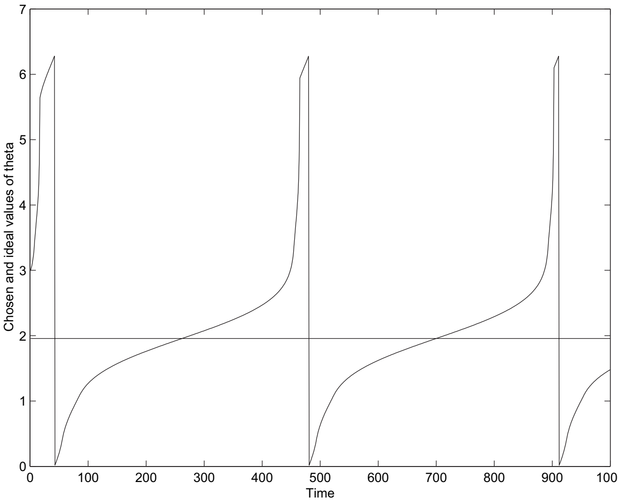

Fig. 3. Variation of θ according to the described control scheme for constant channel and 20 file transfer users with $B = 10000$ and $\gamma * E[SNR] = 4$.

amount of time near optimum. The problem is that the optimal value of θ keeps changing from slot to slot because of the time-varying nature of the fading channels. Thus a search around the phase space seems necessary. Consider now the two parameters B and γ of the scheme as given in (12). A higher value of γ means that the scheme will result in θ spending a larger fraction of time at values where $SNR^*_l(\theta(l))$ is large. But large values of γ will lead to slow changes in the value ofθ, which may prevent it from effectively tracking the optimum. Multiplication by B mitigates this effect somewhat, by ensuring a faster update. Choosing good values of B and γ is a joint problem, and is best done by offline simulations searching over the joint space.

For the remainder of this section we illustrate the working of the control scheme by considering sample paths of a system with 20 file transfer users having Rayleigh channels and a base station with two antennae.

Consider first the case of slow fading channels, for which the channel is constant within the time scale of interest, so that $h_{i,k}(l) \equiv h_{i,k}$. The control scheme is not explicitly designed for this case, but it illustrates the dynamics of the control. In this case the value θ^* of the phase given by $\theta^* = arg(h_{1,k^*} h^*_{2,k^*})$, where $k^* = arg \max_k |h_{1,k}|^2 + |h_{2,k}|^2$, is optimal for all time. In this scenario the trajectory $\theta(\cdot)$ varies more slowly for those l when $|\theta(l) - \theta^*|$ is small (but the variation is always in one direction), as shown in Figure 3.

The ratio of time spent in an ϵ neighborhood of a "good" value of θ to that spent around a "bad" value is approximately $\frac{e^{\gamma SNR_{good}}}{e^{\gamma SNR_{bad}}}$, so it is desirable to have a large value of γ. But if γ is too large, it takes a long time to search through the state space, and θ may spend long periods in bad states, which is not desirable in case the channel is not constant. One option is to choose the corresponding value of B so as to avoid this.

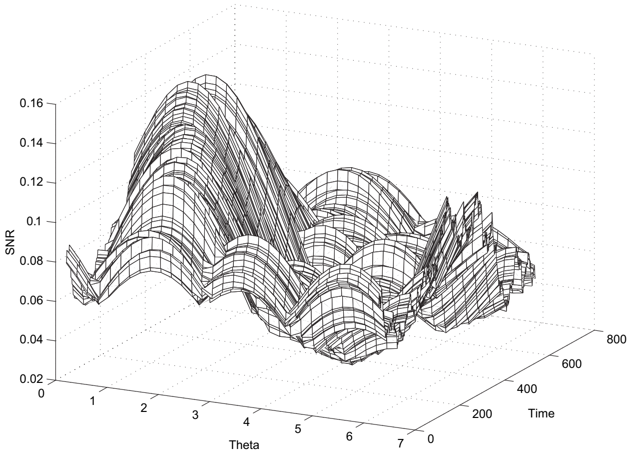

Fig. 4. A typical reward function $Y(l, \theta)$ for a A = 0.999 and 20 users.

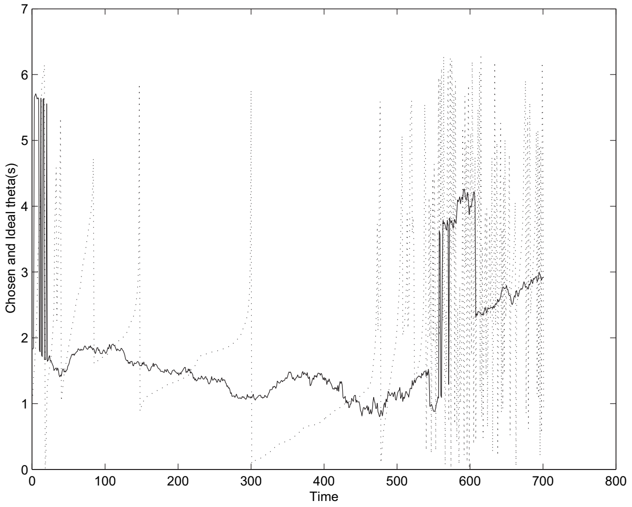

Fig. 5. Ideal $\theta^*(l)$ and controlled $\theta(l)$ for the $Y(\theta, l)$ of figure 4.

Now consider the same scenario with fast fading channels - users have Rayleigh channels whose evolution in time is described by (2). Consider the reward function Y defined by

$$Y(l,\theta) \quad \triangleq \quad \max_k \frac{1}{2} \left| h_{1,k}(l) + e^{j\theta} h_{2,k}(l) \right|^2$$

which is shown in Figure 4 for a particular sample path. For the low SNR regime the objective function f to be maximized, given by (7), can be restated in terms of Y as

$$f(\theta(\cdot)) \quad \triangleq \quad \frac{1}{L} \sum_{l=1}^{L} Y(l, \theta(l))$$

and the optimal trajectory θ^* which maximizes this is given by

$$\theta^*(l) \quad \triangleq \quad arg \left(h_{1,k^*_l}(l) h^*_{2,k^*_l}(l) \right)$$

where

$$k^*_l \quad \triangleq \quad arg \max_k \frac{1}{2} \left(|h_{1,k}(l)|^2 + |h_{2,k}(l)|^2 \right) \qquad (13)$$

Figure 5 shows $\theta^*(\cdot)$ and $\theta(\cdot)$ varied according to our control strategy. We can see it "tracking" $\theta^*(\cdot)$ in the sense of slowing down when $|\theta(l) - \theta^*(l)|$ is small.

Figures 4 and 5 also illustrate another interesting aspect of the scheme: during times when the channel is "good" (i.e. when $Y(l, \theta^*(l))$ is relatively large) the tracking is good (e.g. in time slots 10 to 500) and when the channel is bad the $\theta(l)$ does not slow down much but rather keeps sweeping the phase space rapidly until the next "good" channel.

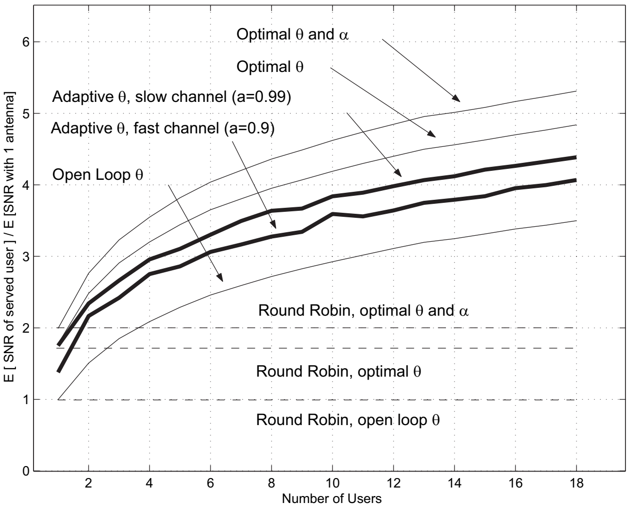

Fig. 6. Performance of the adaptive scheme for θ variation for Rayleigh fading in the low SNR regime.

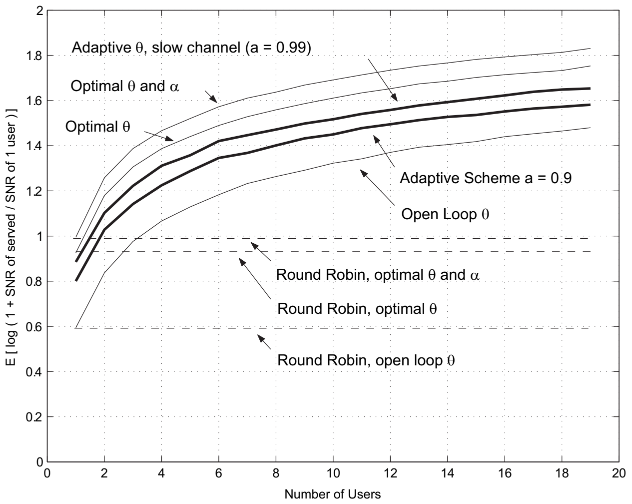

Fig. 7. Performance of the adaptive scheme for θ variation for Rayleigh fading in the high SNR regime.

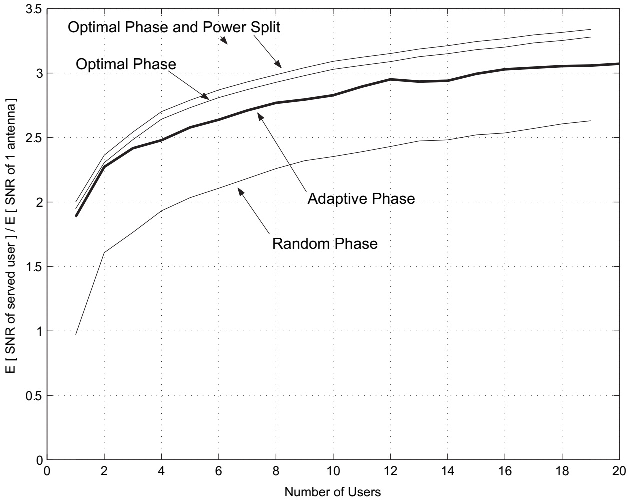

Fig. 8. Performance of the adaptive scheme for θ variation for Ricean fading in the low SNR regime.

V. RESULTS 1: FILE TRANSFER USERS

In this section we present the performance of the proposed scheme for the scenario when all users are of file transfer type. The users are assumed to be symmetric, i.e. all users have independent channels that have the same type, strength and channel speed. In the case of Rayleigh fading the channels follow the update (2) and in the case of Ricean fading they follow (4). In both cases the values of σ^2 and a are the same for all users.

Section II-E formulates the problem formally for this scenario. For the low SNR regime (7) gives the average throughput for a particular scheme π, while (8) does the same for the high SNR regime. Recall that for both SNR regimes and both kinds of channels, the throughput f for schemes corresponding to the benchmark and bounds presented in Section III do not depend on the speed of variation a of the channels, but just on their marginal distribution at any one time.

For the case of Rayleigh fading in the low SNR regime, Figure 6 plots this average throughput for the following cases:

1) When θ is varied in an open loop fashion (see Section III-A). This is denoted in the figure as "Open Loop θ". As mentioned in Section III-A, this is the performance independent of the number of antennas, as long as θ is not controlled using feedback. In particular, it is the best performance one can obtain if there is only one antenna at the base station.
2) When θ is chosen optimally given perfect channel information (see Section III-B) for two antennas at the base station. This is denoted in the figure as "Optimal θ". This requires perfect feedback of the phase of each of the channels to each user.
3) When θ is chosen optimally *and* the power split α between the two BS antennas is also chosen optimally (see Section III-B). This is denoted by "Optimal θ and α". This requires perfect feedback of the phase and amplitudes of each channel to each user.
4) When the power is split equally and θ is varied according to the scheme given by (12), for channel speeds $a = 0.9$ (fast) and $a = 0.99$(slow) in the terminology of (2). These are denoted by "Adaptive θ".
5) And finally, as a comparison, we also plot using dashed lines the performance when pure round-robin scheduling is used.

Figure 6 shows that there are significant gains to be had from exploiting multiuser diversity. Adding on the adaptivity in phase using the scheme (12) proposed in Section IV gives further increases in SNR using very minimal feedback and a simple control. Exact phase and amplitude feedback (and the resulting transmission strategy) could result in further gains, but achieving these gains takes more work than the proposed adaptive scheme.

The average throughput for the cases listed above in the high SNR regime is shown in Figure 7.

The performance of the scheme as compared to the benchmark and upper bounds is presented in Figure 8 for the low SNR regime. For the purposes of simulation and scheme evaluation, in our model for Ricean fading we assumed a Rice factor κ of 8.

Figures 6, 7, and 8 illustrate an important observation about the performance of the scheme: it depends on the channel speed. Slower channels result in higher throughput, as can be seen by the throughput curves for speeds of $a = 0.9$ (faster channel) and $a = 0.99$ (slower channel). A slower channel speed means that a strong user remains strong longer, and so the speed-varying trajectory through the phase space can spend a larger fraction of its time closer to its optimal value. An extreme case that serves to illustrate this phenomenon is when the channel is fixed over the time scale of interest, as was shown in Figure 3. Here the trajectory for θ under the adaptive scheme is seen to spend a large fraction of its time "near" the optimal value θ^* of the system (which remains constant with time since the channels do not evolve). In fact, as the coherence time becomes large, corresponding to $a \to 1$, choosing a very large value of γ and some fixed value of B will result in the performance of the adaptive scheme converging to the upper bound that uses perfect phase information.

The appendix contains simulations for Rayleigh fading in the low SNR regime for file transfer users but channels evolve in time according to the Jake's model. The observations made in the paragraph above continue to be the case for Jake-s model as well.

It has to be mentioned here that the scheme throughput values shown in all the scenarios above were achieved by finding parameters B and γ of the adaptive scheme given in (12) for each scenario by a search through the joint space of (B, γ). Such a search is time consuming, but it can be done off-line and the "optimal" choices for each regime and number of users scenario can be stored.

A natural question that arises is the effect of deployment of such an induced fluctuation scheme in a cellular context, and the effect of out-of-cell interference. This question is briefly addressed in [1] for the case of randomly induced fluctuations. Research on the effects of induced fluctuations in a cellular context is presently being carried out by Dangui and Hajek [8]. Initial simulation results for the case of two interfering cells suggest that inter-cell interference is burstier then thermal noise (due to Rayleigh fading) and so the efficiency of harnessing multiuser diversity is better than predicted by using thermal noise as the only interference source. Also, whether or not the interfering base station is inducing fluctuations, adaptively or otherwise, makes little difference.

VI. RESULTS 2: REAL TIME USERS

This section presents the performance of the scheme when all the users are of the real time type, as defined in Section II-D. Besides the basic symmetric scenario where all users having similar channels, we will also consider the asymmetric scenario where a portion of the users have statistically stronger channels than others.

Fairness is an important issue for users with queues that are sensitive to delays in service (as opposed to file transfer users). If the channels are statistically identical, the proposed scheme will give the same average delay for each user, when considered over a long enough time horizon. Often, however, there is a need for fairness over short time scales. Many explicit fairness mechanisms have been proposed in this regard for systems that involve communication over time-varying channels to multiple users - for example the exponential rule by Shakkottai and Stolyar [5]. In this section we will investigate the effect of having one such explicit fairness mechanism on the delay and throughput of the system, and also investigate how it interacts with the proposed adaptive scheme to control θ.

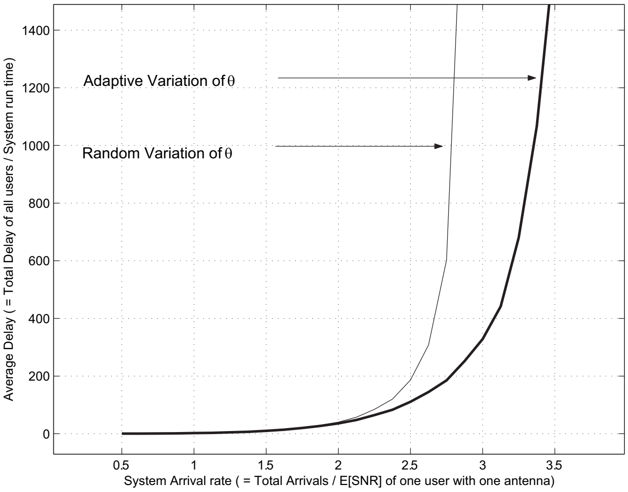

Fig. 9. Average delay vs. system throughput for symmetric users - Effect of controlling phase.

All the results presented in this section are for Rayleigh fading in the low SNR regime, with the channel speed parameter $a = 0.999$, which corresponds (approximately) to a coherence time of a thousand slots.

A. Symmetric Channels

The performance of the scheme to control θ, given by (12), is compared to a scheme that at each instant picks a value of θ that is random and uniformly distributed in $[0, 2\pi]$, and independent of past values or channel feedback. As mentioned before, in both cases the user selected for service is the one that has the maximum potential amount of data to transmit. The average delay in the system for each scheme is given by (9).

For 10 users with Rayleigh fading channels in the low SNR regime, the values of average delay for each of the two schemes above are plotted in Figure 9 as a function of the system arrival rate Kmp.

It is appropriate at this juncture to comment on a tradeoff operating between the open loop and adaptive schemes. On one hand, a scheme for the adaptive variation of θ that "sticks around with the strong user" serves that specific user extremely well at the expense of the other users in the system, who see longer delays until their channels become favorable. Random θ variation on the other hand cycles through the users more quickly but serves less data each time. It is not a-priori clear in which scenario the delays will be lower. The results in Figure 9 show that the benefit of adaptive control outweighs the effect of sticking to the same user for a long time.

It is interesting to observe the effect of adding fairness mechanisms in the scenario when all users have statistically

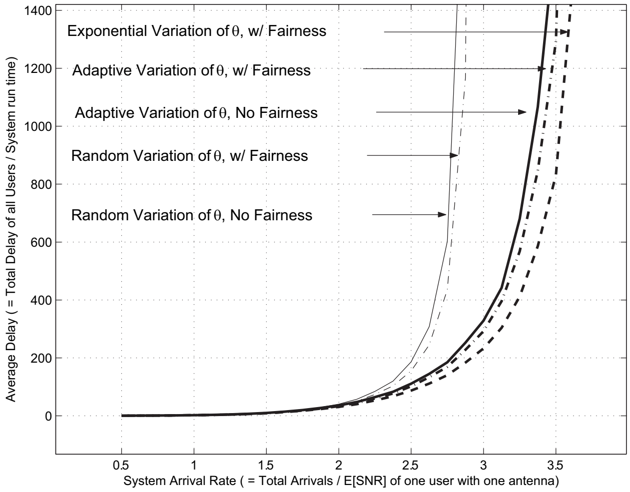

Fig. 10. Average delay vs. system throughput for symmetric users - Effect of fairness mechanisms.

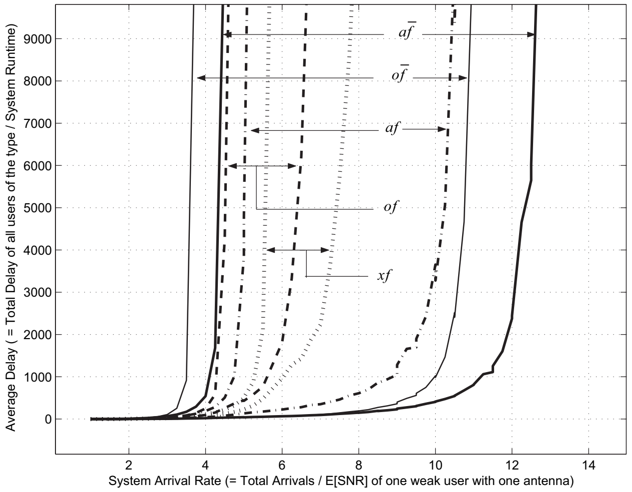

Fig. 11. Average delay vs. system throughput for asymmetric users.

similar channels. As mentioned, statistical symmetry will result in all the users seeing the same long-run average delay, so it is already "fair" in that sense. The addition of an explicit fairness mechanism can have two effects on the service:

E1 On the one hand, it decreases the amount of service that occurs in a *particular* slot for the same channels given that all users have the same amount of data to transmit in that time slot.

E2 On the other hand, it distributes the load on the system over time, so that there will be slots where users that otherwise would have been empty or low on data to send now have data and so can compete. The amount of service in these slots is more than in the absence of the fairness mechanism.

For the purposes of our investigation we adopted the *exponential rule* ([5], see Appendix B) for our simulations. Figure 10 shows the performance for five schemes - four of which are combinations of random and controlled (according to the scheme in Section IV) phase variation and the presence and absence of the exponential fairness mechanism. For both the random and adaptive variation of θ it is seen that effect E2 dominates, leading to improvements in the performance when the fairness mechanism is added.

Another idea is to modify the update rule itself, so that it is matched not to the SNR of the best user but to the index b of the exponential rule, as mentioned in Appendix B. This is the fifth scheme, denoted as "Exponential Variation" in Figure 10. Instead of the update rule given by (12), exponential variation means that the update rule is

$$\begin{aligned}\Delta\theta(l+1) &\overset{\triangle}{=} \theta(l+1)-\theta(l)\\ &= Be^{-\gamma b^*(l)}\end{aligned} \tag{14}$$

where $b^*(l) = \max_i b_i(l)$ and the $b_i(l)$ are the indices of the exponential rule, as defined in Appendix B. Service is again given to the user with the highest index b_i. It is seen that this scheme gives a further improvement over the original adaptive scheme. Apparently, if it is deemed best to serve a particular user, taking both signal strength and queue size into account, then it helps to use the same considerations for the control of θ. This would avoid the eventuality of service being given to one user while the θ "adapts to" another user.

B. Asymmetric Channels

We also considered the case when users do not have the same channel strengths, i.e., the channels for some users are statistically stronger than those of other users. Here the issue of fairness arises naturally:

1) Channel asymmetry means that even in the long run, the experienced delay may vary significantly across different kinds of users in the absence of any fairness mechanism.
2) For certain arrival rates, a fairness mechanism would allow all users to have stable queues, while the absence of such a mechanism may lead to unstable queues for statistically weaker users.

Thus a fairness mechanism can contribute not only to lessening the disparities in the delays but also to increasing the system throughput (in the sense of the maximum possible arrival rate it can support while keeping all users stable).

We considered the case of a system of 20 users with Rayleigh channels, 10 of whose channels are twice as strong as those of the remaining 10. Referring to the channel model described in Section II-B, and specifically the channel update Equation (2), if the 10 weaker users have $w_{i,k}(l) \sim \mathcal{CN}(0, 2\sigma^2)$ (and hence $h_{i,k}(l) \sim \mathcal{CN}(0, 2\sigma^2)$), the 10 stronger users have $w_{i,k}(l) \sim \mathcal{CN}(0, 4\sigma^2)$ (and hence $h_{i,k}(l) \sim \mathcal{CN}(0, 4\sigma^2)$).

The plots of delay vs. throughput for various schemes are given in Figure 11. The nomenclature in the labeling is:

a = Adaptive θ f = with fairness mechanism
o = Random θ $\overline{f}$ = without fairness mechanism
x = Exponential Variation

Each curve in the figure is the performance of one type of scheme for controlling θ, in the presence or absence of fairness, for one kind of user. Thus for example the label $o\overline{f}$ corresponds to θ being randomly varied from slot to slot and service being given to the user with the highest potential amount of data to send, irrespective of their queue lengths (i.e. there is no fairness mechanism). The two curves it points to correspond to the delays for the stronger and weaker users.

For all the labels, the curve on the left plots the delay seen by the weaker users and the one on the right plots the delay for the stronger users.

The following observations can be made about the performance of the scheme (with references to the corresponding lines in the figure):

1) In the absence of a fairness mechanism, adaptation of θ lowers the average delay of all users. (Compare $o\overline{f}$ to $a\overline{f}$)
2) In the presence of a fairness mechanism, the adaptation of θ leads to lower delays for all users, with the gain for stronger users being much larger than that for the weaker users, as compared to random θ variation. (Compare of to af)
3) Adding the fairness mechanism significantly shrinks the gap in service between the strong and weak users, both for adaptive and random θ variation. (Compare $o\overline{f}$ to of and $a\overline{f}$ to af)
4) As mentioned in Section VI-A, the adaptive scheme may adapt to the strongest user while the fairness mechanism may force service to another user - and this eventuality can be avoided by exponential variation. When the users are asymmetric, this mismatch in the update mechanism of the original scheme leads to higher delays for the weaker users as compared to the exponential variation scheme. In fact, of all schemes mentioned in this section, it is seen that the exponential variation scheme gives the best performance for weaker users. This of course comes at the expense of the stronger users. (Compare of, af and xf)

VII. CONCLUSIONS

This paper proposes a scheme for and shows the advantages of inducing adaptive fluctuations in a cellular wireless system utilizing multiuser diversity and having channels with memory, in a range of different scenarios:

1) Users having infinite backlog or finite queues.
2) The performance measures of interest being throughput or delay.
3) The channels having Rayleigh or Ricean fading.
4) The channels of the users being symmetric or asymmetric.
5) System operating in the high or low SNR regime.
6) Presence or absence of additional fairness mechanisms.

Adaptively inducing fluctuations leads to better performance for all users in all scenarios, in comparison to when there are no induced fluctuations or when the induced fluctuations are not adapted.

For infinitely backlogged users communicating over channels varying in time but with memory, it is seen that the average rate of service is much higher for adaptively induced fluctuations than when there are no induced fluctuations or when the fluctuations are not induced adaptively. The performance is also compared to upper bounds derived from beamforming (which are unachievable given only SNR feedback), and it is seen that the more the memory of the channel, the closer the performance of the scheme is to the upper bound.

For real time users with finite queues fed by a Bernoulli arrival process and again communicating on channels with memory, adaptation of induced fluctuations led to lower average delay (averaged over the users). The working of the scheme in conjunction with the exponential rule for fairness was also investigated, both for when users are identical and in the presence of asymmetry between the users (i.e. one set of users having statistically stronger channels than another set). Adaptation leads to lower delays in scenarios with and without an additional fairness mechanism for all users. Further, in the presence of a fairness mechanism, it is seen that matching the adaptive scheme to the index of the mechanism instead of to the highest SNR leads to better performance.

Future work could involve investigating the case of more than two antennae at the base station and the effects of other cell interference.

APPENDIX I
THE MAXIMUM OF SUMS OF EXPONENTIALS

In this appendix we present a method for finding the expected value of the right hand side of (10) in Section III-B. Since $|h_{i,k}|^2$ has the exponential distribution with mean one (exp(1)) for each i, k, the problem is equivalent to the problem of finding the maximum of K independent random variables, each of which is the sum of N independent $exp(1)$ random variables. We calculate it explicitly for the cases when $N = 1$ and $N = 2$. We also show that in the asymptote of large K, the distribution of the maximum for $N = 2$ belongs to a particular extremal type. In arriving at this extremal type formulation we use a scaling that gives us a good approximation for the expected value for finite K as well.

A. Explanation of the Method

Consider NK (N and K are positive integers) independent random variables $\{h_{n,k}\}_{n=1,k=1}^{n=N,k=K}$ each of which is exponentially distributed with parameter 1. The quantity we are interested in calculating is $E[\max_k(h_{1,k} + \ldots + h_{N,k})]$.

The N^{th} jump of a Poisson process is distributed like the sum of N independent exponential random variables, with parameter 1. Consider K independent Poisson processes and define $r_i(t)$ to be the number of processes that have undergone i or more jumps by time t. Note that for any $i < j$, $r_i(t) \geq r_j(t)$. Then, the N dimensional random process $\underline{\mathbf{r}}(t) \overset{\triangle}{=} (r_1(t), \ldots, r_N(t))$ is Markov. The quantity we are interested in, namely, $E[\max_k(h_{1,k} + \ldots + h_{N,k})]$, is the expected time needed for $\underline{\mathbf{r}}(t)$ to hit state $(K, \ldots, K)$ starting from $(0, \ldots, 0)$, i.e., $E[\min\{t : \underline{\mathbf{r}}(t) = (K, \ldots, K)\}]$.

This expected time can be calculated recursively using the transition probabilities of the states. We give a more detailed presentation of the method for $N = 1$ and $N = 2$, since these are the cases of primary interest.

B. N = 1

Consider $N = 1$. We have to find $E[\min\{t : r_1(t) = K\}]$. The chain is one-dimensional and the transition rates are as shown in Figure 12. From this it is seen that, for all K,

$$E[\max_k h_{1,k}] = 1 + \frac{1}{2} + \ldots + \frac{1}{K}$$

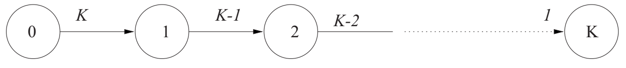

Fig. 12. Markov chain used for evaluating expected value of the maximum of K exponentials - the expected value corresponds to the expected hitting time of state K from 0.

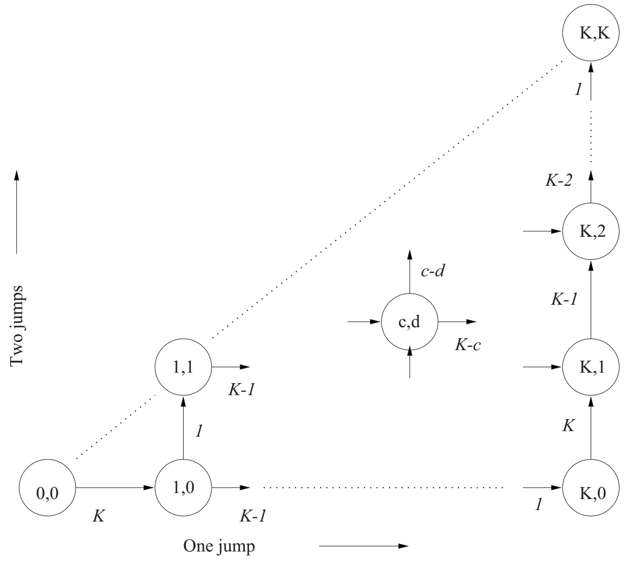

Fig. 13. Markov Chain used for evaluating expected value of the max of K sums of two exponentials - the expected value corresponds to the expected hitting time of state (K, K) from $(0, 0)$.

$$Var[\max_k h_{1,k}] \;=\; 1+\frac{1}{2^2}+\ldots+\frac{1}{K^2} \;<\; \frac{\pi^2}{6}$$

C. N = 2

For the case of $N = 2$, in addition to the two-dimensional Markov chain used for the exact calculation we also show that as $K \rightarrow \infty$ the distribution of the maximum belongs to a certain extremal type. The two-dimensional Markov chain used for the evaluation is shown in Figure 13. Let $s_{c,d}$ denote the expected amount of time to hit state (K, K) when the process starts from state (c, d). We have to evaluate $s_{0,0}$. From the transition rate diagram it can be seen that

$$s_{c,d} \;=\; \frac{1}{K-d}+\frac{c-d}{K-d}s_{c,d+1}+\frac{K-c}{K-d}s_{c+1,d} \quad (15)$$

for $c = 0, \ldots, K$ and by definition $d = 0, \ldots, K-1$. By definition, $s_{K,K} = 0$. We can thus use this update equation to calculate $s_{0,0}$ starting from $s_{K,K}$.

The extremal types theorem (Theorem 1.4.2, p. 11, [6]) gives the convergence in distribution of the maximum of i.i.d. random variables under appropriate scalings and shiftings. It can be shown that the cumulative distribution function of M_n satisfies

$$F_{M_n}(x+\log n+\log\log n) \quad\rightarrow\quad exp(-e^{-x})$$

This suggests the approximation $E[M_n] \approx \log n + \log\log n$ (see [6]). A little bit of tweaking of this function shows that $\max_{2\leq n\leq 1000} |E[M_n] - log(n+1) - log(1+log(n)) - 1| \leq 0.1248$, giving a good approximation for the maximum of n terms for $n \leq 1000$. The two curves $E[M_n]$ and $\log(n+1)+\log(1+\log n)+1$ are plotted in Figure 14.

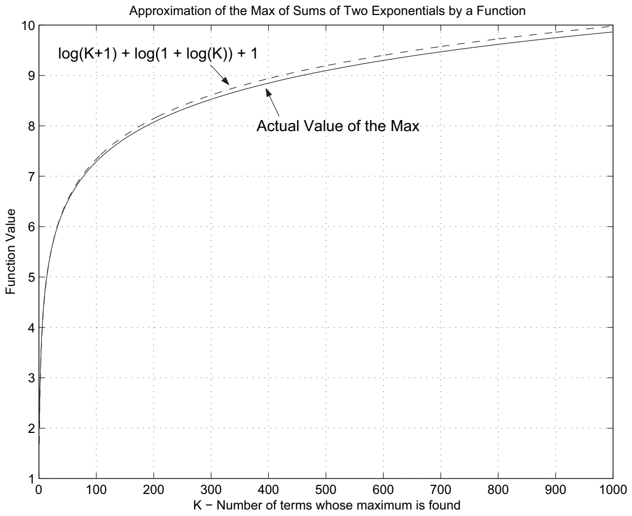

Fig. 14. Approximation of the max of sums of two exponentials by $\log(n+1)+\log(1+\log n)+1$.

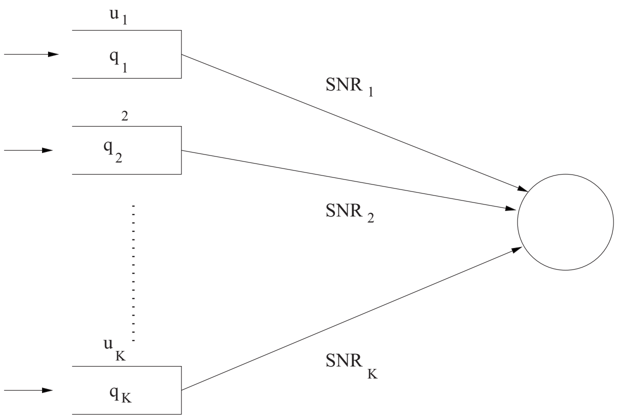

Fig. 15. A system with server, queued users, and time varying channels.

APPENDIX II
THE EXPONENTIAL RULE

For a system consisting of one server serving many queues over time-varying channels, Shakkottai and Stolyar [5] proposed the *Exponential Rule* for the fair scheduling of service. The following is a brief explanation of the rule in the context of this paper. More specifically, it explains the fairness mechanism we have investigated in conjunction with the adaptive scheme in Section VI.

Consider a system like the one shown in Figure 15 consisting of a base station serving K users $\{u_k\}_1^K$ over time varying channels with SNRs $\{SNR_k(t)\}_1^K$. At any time t, users have some data $\{q_k(t)\}_1^K$ to send in their respective queues. Assume the channels are ergodic and have mean SNRs $\{\overline{SNR}_k\}_1^K$, respectively. Assume also that scheduling a user k at a particular time t means that at most $SNR_k(t)$ of its queue can be emptied by the server.

At each time, assign an index $b_k(t)$ to each user, given by

$$b_i(t) \;\triangleq\; \frac{SNR_k(t)}{\overline{SNR}_k} exp\left(\frac{a q_k(t)}{1+\sqrt{\frac{1}{K}\sum_{k=1}^{K} q_k(t)}}\right)$$

Then, the exponential rule states that at time t the server

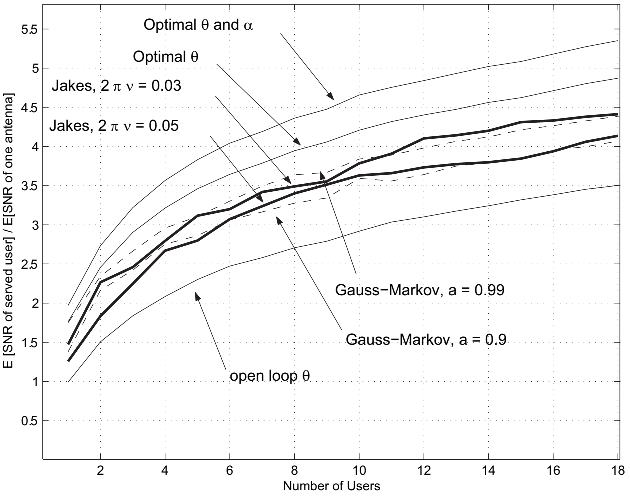

Fig. 16. Performance of the adaptive scheme for θ variation for Rayleigh fading in the low SNR regime with Jake's model.

should pick for service the user with the highest value of the index, i.e., $\arg\max_{1\leq k\leq K} b_k(t)$.

For the purposes of the simulations of this paper we took the parameter a of the exponential rule to be 1 for all time.

Other fairness rules have been proposed that can be potentially useful in a scenario like the one considered in this paper, see for example [4].

APPENDIX III
SIMULATIONS FOR JAKE'S MODEL

The simulations presented in the main body of the paper were generated using a Gauss-Markov fading model. In this appendix we report simulations using the Jake's channel model [9]. For the purposes of our simulations we used the model for Rayleigh fading described below, taken from [10].

We assume the channel from each antenna at the base station to the antenna at each mobile is independent of all other channels, and that all channels are identically distributed. The evolution in time of one of these channels, denoted here by $h(t)$, is given by

$$h(t) \;=\; \sqrt{\frac{1}{P}}\sum_{p=1}^{P}\exp\left(j\left(\psi_p - 2\pi\nu_{max}t\cos\frac{2\pi p}{P}\right)\right)$$

Here P is the number of paths, ψ_p is the phase of path p, $\frac{2\pi p}{P}$ is the angle of arrival of path p and ν_{max} is the maximum frequency. The phases ψ_p are chosen to be i.i.d. uniform in $[0, 2\pi]$. Since the path angles are regularly spaced and the paths have equal power, the model corresponds to *isotropic* Rayleigh fading.

The rest of the system model and problem statement remains unchanged from the one described in Section II for file transfer users in the low SNR regime.

By the central limit theorem the marginal distribution of the channel at any instant in time is nearly the same as in the Gauss-Markov case. Hence the performance when optimal beamforming is used, either only in phase or in both phase and amplitude, will remain as before. Similarly, the performance of open loop control strategies will also be the same as in the Gauss-Markov fading model.

Figure 16 presents simulation results on the performance of the adaptive scheme proposed in this paper for two scenarios: a faster channel with $\nu_{max} = 0.05$ and a slower one with $\nu_{max} = 0.03$. These correspond, approximately, to cyclic channel variations with periods of 60 and 100 time slots respectively. As noted in the Gauss-Markov case, the scheme performs better when the channel varies more slowly, and adaptive induced variation yields larger throughput than open loop induced variations. The simulations show that the Jakes model and Gauss-Markov model yield qualitatively and quantitatively similar predictions about the relative performance of adaptively induced fluctuations. The curves "optimal θ", "optimal θ and α" and "open loop θ" are the same as in the Gauss-Markov simulations of figure 6. The labeling in figure 16 is similar to that of Figure 6. Also plotted for comparison in figure 16 are the performance curves for the scheme under Gauss-Markov fading.

REFERENCES

[1] P. Viswanath, D. Tse, and R. Laroia, "Opportunistic beamforming using dumb antennas," *IEEE Trans. Inform. Theory*, vol. 48, no. 6, pp. 1277-1294, June 2002.

[2] S. M. Meerkov, "Deceleration in search of the global extremum of a function," *Avtomatika i Telemekhanika*, no. 12, pp. 129-139, Dec. 1972.

[3] F. Nesser and J. Massey, "Proper complex random processes with applications to information theory," *IEEE Trans. Inform. Theory* vol. 29, no. 4, pp. 1293-1302, July 1999

[4] X. Liu, E. K. P. Chong, and N. B. Shroff, "Opportunistic transmission scheduling with resource-sharing constraints in wireless networks," *IEEE J. Select. Areas Commun.*, vol. 19, no. 10, pp 2053-64, Oct. 2001.

[5] S. Shakkottai and A. Stolyar, "Scheduling for multiple flows sharing a time-varying channel: the exponential rule," *American Mathematical Society Translations, Series 2*; a volume in memory of F. Karpelevich, Y. M. Suhov, Editor, vol. 207, 2002.

[6] M. R. Leadbetter, G. Lindgren, and H. Rootzen, *Extremes and Related Properties of Random Sequences and Processes*, 1983.

[7] H. Holma and A. Toskala, *WCDMA for UMTS: Radio Access for Third Generation Mobile Communications.* Wiley, 2000.

[8] R. Dangui, M.S. thesis, under preparation.

[9] W. C. Jakes, *Microwave Mobile Communications*. Wiley, 1974; reprinted by IEEE Press in 1994.

[10] V. V. Veeravalli and A. M. Sayeed, *Wideband Wireless Channels: Statistical Modelling, Analysis and Simulation*. Book in preparation for publication by John Wiley and Sons.

Sujay Sanghavi received a B. Tech. in Electrical Eningeering from the Indian Institute of Technology, Bombay in 2000 and an MS in Electrical and Computer Engineering from the University of Illinois at Urbana Champaign in 2002.

Since 2002 he has been pursuing his PhD in ECE at UIUC. His research and academic interests include communication systems, random networks, wireless networks and game theory.

Bruce Hajek (M'79-SM'84-F'89) received a B.S. in Mathematics and an M.S. in Electrical Engineering from the University of Illinois at Urbana-Champaign in 1976 and 1977, and a Ph. D. in Electrical Engineering from the University of California at Berkeley in 1979. He is a Professor in the Department of Electrical and Computer Engineering and in the Coordinated Science Laboratory at the University of Illinois at Urbana-Champaign, where he has been since 1979. His research interests include communication and computer networks, stochastic systems, combinatorial and nonlinear optimization, and information theory. He served as Associate Editor for Communication Networks and Computer Networks for the *IEEE Transactions on Information Theory* (1985-1988), as Editor-in-Chief of the same Transactions (1989-1992), and as President of the IEEE Information Theory Society (1995).

Dr. Hajek was a winner of the 1973 USA Mathematical Olympiad, and he received the Eckman Award of the American Automatic Control Council (1982), an NSF Presidential Young Investigators Award(1984), an Outstanding Paper Award from the IEEE Control Systems Society (1985). He is a member of the US National Academy of Engineering and he received the 2003 IEEE Kobayashi Computer Communications Award.