

Optimal Link Adaptation in Wideband CDMA Systems

Tim Holliday, Andrea Goldsmith, and Peter Glynn
Stanford University

Abstract—We develop a general framework for optimizing link adaptation for multiuser CDMA systems in the wideband limit. The framework is then used to solve for the optimal power control policy that minimizes average transmit power while satisfying a constraint on the per-user probability of packet loss due to deadline expiration. The optimal link adaptation is found through an infinite horizon dynamic program. Typical dynamic programming formulations do not perform well for CDMA systems since the size of the problem grows exponentially with the number of users. We show that in the limiting regime of long spreading codes and large numbers of users the problem size collapses to that of a single user formulation, allowing us to solve previously intractable problems. In particular, we consider a concrete example of power control in a CDMA system with deadline constrained traffic. We solve for the optimal power control policy and examine the tradeoffs between power consumption, probability of deadline expiration, and number of users in the system. Finally we present simulation results evaluating the accuracy of the wideband limit when used as an approximation for finite bandwidth systems. We show that the optimal power control and resulting performance for the limiting regime is a reasonable approximation for the large bandwidths expected in next generation wireless systems.

I. INTRODUCTION

Next generation wireless systems promise the introduction of a wide range of multimedia services. Streaming audio/video, voice over IP, and other services will share the network with traditional voice and data traffic. These new services require significant guarantees on system performance, including constraints on throughput, delay, voice/video quality, and probability of data loss. Wireless links are inherently unreliable since the transmission medium is subject to time-varying fading and interference. In order to combat adverse propagation conditions mobile radios can adapt various controls (e.g. transmission power, source/channel codes, etc.) to the wireless channel and performance requirements.

In this context, there has been a great deal of research on link adaptation policies for delay constrained traffic [1], [2], [5], [6]. Most of the literature in this area focuses on adapting power and source/channel coding to minimize average power consumption subject to performance constraints. As well, nearly all of these formulations make a fundamental simplifying assumption about the nature of the interference in a wireless system. That is, the interference observed by any user in the system does not react to the adaptive control chosen by the user. This simplification allows the interference power to be modeled as a fixed constant or independent random process. The interference model then permits an optimal link adaptation formulation that assumes each mobile user operates in isolation and does not impact the performance of other users in the system. Typically authors attempt to justify this assumption in time-division systems by the relatively large distance separating mobile devices that occupy the same time slot. While this might be a reasonable assumption in TDMA systems, it is not sufficient for code division (CDMA) systems where users occupying the same time and frequency slots reside in the same cell. Indeed, the power control required to combat the “Near-Far” problem in CDMA clearly shows that the interference in a code division system cannot be adequately modeled as unresponsive.

Previous research for delay constrained link adaptation, and power control in particular, does not easily extend to account for

responsive interference. Literature in this area consists of dynamic programming or other stochastic control solutions that allow for a very general modeling framework and a wide variety of performance constraints. However, when the assumption of unresponsive interference is removed, the size and complexity of these formulations typically grows exponentially with the number of users in the system. As a consequence, the well known “curse of dimensionality” prevents solution of the optimal link adaptation strategy for all but the most trivial models of multi-user CDMA systems. Obviously this issue is only magnified as we consider wideband systems with larger numbers of wireless devices.

In this paper we present a dynamic programming formulation for optimal link adaptation in CDMA systems that does not suffer from the dimensionality problem mentioned above. In order to accomplish this goal we propose a mean field approximation for the interference in a CDMA cell. Specifically, we consider the evolution dynamics of a CDMA cell in a limiting regime with long spreading sequences and large numbers of users. We will show that within this regime the size of the optimal link adaptation problem for a K -user system collapses to that of a single user system with the addition of non-linear dynamics. This reduction allows us to construct and solve a general class of delay constrained optimal link adaptation problems for CDMA systems with responsive interference that were previously intractable. We illustrate our technique by finding the the optimal power control policy that minimizes average transmit power subject to a delay constraint.

Several authors have proposed mean field approximations for the signal to interference ratio (SIR) in CDMA systems for a variety of receiver structures [11] and system models. In this area of research the approximation considers the limiting SIR as the number of wireless devices (K) and spreading sequence lengths (N) tend to infinity while the ratio $\frac{K}{N}$ remains fixed. In this paper we will consider the same limiting regime. However, rather than approximate the SIR in a single timeslot we find an approximation for the evolution equations of the interacting users in a CDMA cell, and use this approximation to solve the optimization problem. Though we find the optimal link adaptation for a system operating in the wideband limit, the system model is only an approximation for a finite user system. Hence our optimal link adaptation for the limiting system is only “approximately optimal” for a finite user system. However, we will present simulation results that show this approximation is reasonable for the numbers of users expected in wideband CDMA systems.

The rest of this paper is organized as follows. In the next section we present the system model and the components required to construct the optimal link adaptation (i.e. dynamic programming) problem. In the third section we develop the mean field approximation and show how it allows us to greatly reduce the complexity of our dynamic program. We present numerical results for optimal power control obtained via the dynamic program as well as simulation results checking the validity of our approximation in the fourth section, followed by our conclusions.

II. SYSTEM MODEL AND PROBLEM FORMULATION

In this section we develop a simple model for a wireless user and a CDMA cell. Our goal is to construct a state space model and evolution equations for each user and then use these components to construct a model for an entire cell. We then use this

Work supported by the Office of Naval Research (under grants N00014-99-1-0698 and N00014-02-1-0003) and the Stanford Networking Research Center.

model to formulate an optimal link adaptation problem. Though our formulation is general, in this paper we will consider a particular example in order to facilitate our presentation. Specifically, we will consider a single CDMA cell consisting of users transmitting deadline sensitive data to a base station. We will solve for the optimal power control assuming each user has knowledge of his channel to the base station, observed interference, and the amount of data queued for transmission in his buffer. Though far more general link adaptation strategies and system models are possible (e.g. see [5], [6]), our goal is to demonstrate the extension of previous research to account for responsive interference rather than a discussion of complicated system models.

A. The Data Traffic and Buffer Model

Consider a group of K wireless users transmitting deadline constrained data to a base station. In each timeslot every user may generate a packet with probability λ . When packet j is generated at time t_j it announces a deadline d_j such that if it is not transmitted by time $e_j = t_j + d_j$ (the extinction time) it is dropped from the data buffer and considered lost. In this paper we will assume all packets have a common deadline $d_j = d$. Therefore we can assume each user has a finite size data buffer capable of storing d time slots worth of data. Let $y_i(t)$ be the state of the i th user's data buffer at time t .

B. The Wireless Channel Model

We assume the variations in the wireless channel gain between each user and the base station can be described by a finite-state discrete-time Markov chain (DTMC). Further assume each user's channel follows an independent copy of the same DTMC. Several authors have proposed Markov models for many different types of wireless channels [4], [8]. The construction of these models follows a common procedure. The range of path gains are divided into a finite number of sections and each section becomes a state in the Markov chain. Then a transition matrix, which determines the probability of jumping from one state to another, is constructed in order to accurately model the channel characteristics (e.g. log-normal shadowing or Rayleigh fading). While we do not require any particular set of channel characteristics, we do require that the jumps of the DTMC occur at the time slot boundaries. We will denote the state of the i th user's channel gain at time t as $z_i(t)$.

C. The State Space Model

We define the state of user i at time t as $x_i(t) = \{y_i(t), z_i(t)\} \in X$, where the set X consists of all combinations of buffer states and channel states for a single user. Using the state of each user we can define a state for the entire system as the empirical distribution $\Pi_K(t)$ over all K users in the cell at time t :

$$\Pi_K(j, t) = \frac{1}{K} \sum_{k=1}^K I_{[x_k(t)=j]}, \quad (1)$$

where $I_{[x_k(t)=j]} = 1$ if mobile k is in state j at time t and zero otherwise (i.e. $I_{[A]}$ is the standard indicator random variable). We should point out that for finite K the empirical distribution $\Pi_K(t)$ is a random vector.

D. Link Adaptation

In each time slot every mobile selects an action $a_i(t) \in A$, where A is a finite set of transmission parameters (e.g. power level, coding scheme, constellation size, and so forth). For simplicity we will restrict our attention to optimal power control though we can certainly accommodate more general link adaptation strategies. Thus $a_i(t)$ is the power level selected by user i at

time t from some finite set of power levels A . We assume each mobile chooses a transmission power according to a power control policy $g(x_i(t), \Pi_K(t))$, which maps the state of the i th user and the empirical distribution $\Pi_K(t)$ into a probability distribution across A . This determines the probability that the i th user takes an action $a_i(t) \in A$ at time t (i.e. we are allowing for randomized controls). We also assume each user is forced to follow the same power control policy. Hence if two mobiles are in the same state at time t they will select their action from the same probability distribution. Finally, we define a cost vector $r(a)$, which determines the cost of choosing action a . In this paper the cost vector r is particularly simple since we are only considering power adaptation (i.e. $r(a) = a$). However, for more general link adaptation strategies the cost can be much more complicated.

E. Receiver Structure and SIR

As in [11], we assume the i th user is assigned a N length spreading sequence S_i , which is constructed as follows:

$$S_i = \frac{1}{\sqrt{N}} \{(V_{i1}, \dots, V_{iN}), \quad i = 1, \dots, K, \quad (2)$$

where the elements V_{ij} are independently and identically distributed with mean zero, variance one, and $E[V_{ij}^4] < \infty$. We also assume the number of users per unit bandwidth converges in the limit (i.e. $\frac{K}{N} \rightarrow \alpha$).

The structure of the receiver at the base station is a linear demodulator. Therefore the received SIR for user i at time t with K users in the system is

$$SIR_K(i, t) = \frac{(c_i^T S_i)^2 a_i(t) z_i(t)}{(c_i^T c_i) \sigma^2 + \sum_{j=1, j \neq i}^K (c_i^T S_j)^2 a_j(t) z_j(t)}, \quad (3)$$

where σ^2 is the background noise power, $a_i(t) z_i(t)$ is the received power from user i , and c_i is the demodulator for user i . Notice that the traditional CDMA approach simply chooses $c_i = S_i$ (i.e. a matched filter receiver), while the optimal MMSE receiver will be much more complicated [9].

The above linear demodulator provides two features that are critical to our analysis. First, for the i th user we can directly compute $SIR_K(i, t)$ if we know the power control policy $g(x_i(t), \Pi_K(t))$, the state of the i th user $x_i(t)$, and the empirical distribution of all users $\Pi_K(t)$. Second, the limiting SIR as the number of users becomes large is simple to compute [11] for many different receiver structures. In our model the j th component of the vector $h(g)$ is the received power from a mobile in state j that is following power control policy g . Since the available power levels are chosen from a finite set and infinite transmission power is not permitted, we know that $h(g) \in [h_{min}, h_{max}]^{|X|}$. Then if the receiver is a matched filter, the SIR for the i th user as $K, N \rightarrow \infty$ with $\frac{K}{N} \rightarrow \alpha$ is simply:

$$SIR_\infty(i, t) = \frac{a_i(t) z_i(t)}{\sigma^2 + \alpha \pi(t) h(g)}. \quad (4)$$

As a technical aside, for this limit to exist we must know [11] that the limiting empirical distribution, $\pi(t)$, exists at time t . That is,

$$\pi(t) = \lim_{K, N \rightarrow \infty} \Pi_K(t) \quad (5)$$

must be a deterministic vector. We will discuss this particular issue in Section III.

F. Evolution Equations

Now that we have defined our state space model, power control policy, and receiver structure, we can set up a system of equations describing the evolution of a CDMA cell over time. We assume the probability of mobile i successfully transmitting a packet of data at time t depends only on $SIR_K(i, t)$. Recall that we also assumed the transitions of the i th user's wireless channel are independent of all other users' states and power choices. Therefore the transition probabilities for the i th user are only a function of $x_i(t)$, the power control $g(x_i(t), \Pi_K(t))$, and the empirical distribution $\Pi_K(t)$. Hence we define a transition matrix $P(g, \Pi_K(t))$ whose jk th element is

$$\begin{aligned} P(g, \Pi_K)_{jk} &= P(x_i(t+1) = k | x_i(t) = j, g, \Pi_K(t)) \quad (6) \\ &= f_k(x_i(t), g, \Pi_K(t)), \quad (7) \end{aligned}$$

where f maps $x_i(t)$, g , and $\Pi_K(t)$ to a probability distribution on the states in X .

As we mentioned above, the empirical distribution $\Pi_K(t)$ is a random vector, and therefore the transition matrix $P(g, \Pi_K(t))$ is also random. As one might guess, the number of different values $\Pi_K(t)$ can take will grow very quickly as the number of users in the system increases. Furthermore, the probability distribution for $\Pi_K(t)$ can be particularly difficult to compute since it involves the summation of K random variables that are neither independent nor identically distributed. Indeed, as we will show in the next section, the random transition matrices resulting from interacting users are precisely the issue that makes this link adaptation problem so difficult.

G. Dynamic Programming Formulation

The term "dynamic program" refers to a wide variety of formulations and optimization techniques used to solve optimal control problems. However, one common feature among all of these techniques is the notion of a value function – a quantity that determines the cost or reward associated with a particular control policy. The goal of any dynamic program is to optimize the value function subject to constraints on system behavior (e.g. average power consumption or probability of deadline expiration). In this paper we use what is known as an infinite horizon value function, which assigns equal weight to the cost associated with the selected control in each time slot. Hence if $R_i(g, t)$ is the random process describing the cost associated with the control for user i at time t , we write the value of the power control policy g as

$$V_K(g, i) = \lim_{T \rightarrow \infty} \frac{1}{T+1} \sum_{t=0}^T E_{x, \Pi} [R_i(g, t)], \quad (8)$$

where the expectation is taken with respect to the state $x_i(t)$ and empirical distribution $\Pi_K(t)$.

In order to solve this dynamic program for the problem formulation discussed in the previous sections we must find the optimal power control g that minimizes our value function (8). At this point we can see how the curse of dimensionality makes a general solution to this problem intractable. In order to optimize (8) we must be able to evaluate the value function for any choice of control. This requires us to either compute the expected cost for each time t or to find a steady-state distribution of $R_i(g, t)$, which then requires knowledge of the distribution of $\Pi_K(t)$. However, as we mentioned above, computing the probability distribution of $\Pi_K(t)$ can be extremely difficult. Indeed it is quite easy to create trivial examples with only three or four users such that computing $\Pi_K(t)$ is computationally infeasible. In the next section we will show that these complexity issues are largely eliminated when we consider the appropriate limiting regime of long spreading sequences and large numbers of users.

III. A MEAN FIELD APPROXIMATION

Our goal in this section is to simplify the evolution equations and value function developed above. In particular, we seek an analytical regime that collapses the random transition matrices (6) to deterministic quantities. The simplified transition dynamics will allow us to substantially simplify our state space description of a CDMA cell and optimize our proposed value function. As mentioned above, the appropriate regime is the limit as $N, K \rightarrow \infty$ and $\frac{N}{K} \rightarrow \alpha$. Then, in order to simplify and optimize our value function (8), we will prove the following propositions in the context of this limiting regime:

- 1) the random transition matrices $P(g, \Pi_K(t))$ converge to deterministic matrices for all t ,
- 2) the evolution equation for the empirical distribution $\Pi_K(t)$ converges to a deterministic non-linear equation,
- 3) this deterministic non-linear equation has a unique fixed point (or steady state),
- 4) the existence of a unique steady state simplifies our value function to a simple inner product of two vectors,
- 5) the simplified value function allows us to formulate a standard optimization problem that minimizes our value function over all possible adaptation policies.

Due to space constraints we will not be able to provide all of the mathematical details in the following proofs. The reader is directed to the journal version [7] of this paper for more details.

A. Transition Matrices and the Empirical Distribution

Assume that the K mobiles' states are initialized at time $t = 0$ independently according to a probability distribution $\Pi_K(0)$. Further assume the following limit exists

$$\lim_{K, N \rightarrow \infty} \Pi_K(0) = \pi(0). \quad (9)$$

Then at time $t = 0$, our transition matrix for a single user converges to a deterministic quantity,

$$\lim_{K, N \rightarrow \infty} P(g, \Pi_K(0)) = P(g, \pi(0)). \quad (10)$$

The reasoning behind this statement is straightforward. From [11] we know the SIR for each user converges as $K, N \rightarrow \infty$ provided the empirical distribution $\pi(0)$ exists. Since our transition matrices for each user are determined by user's SIR and control policy we should expect the matrices to converge. In order to see how this affects our transition dynamics for time $t > 0$ we consider the following proposition:

Proposition 1: As $K, N \rightarrow \infty$ the interference at time $t = 1$ converges to the deterministic non-linear matrix function,

$$\pi(1) = \pi(0)P(g, \pi(0)), \quad (11)$$

and by induction this equation holds for general time t .

Proof: From the result in (10) we know that each of the mobiles in the system will transition from time $t = 0$ to time $t = 1$ using the same transition matrix $P(g, \pi(0))$. That is, for a given power control g and empirical distribution $\pi(0)$ each of the users in the system will transition to the next time slot independently of all other users. Hence, if we group the users in the system based on their state at time $t = 0$, we can view the next state transitions of each of these groups as independent multinomial random variables. Specifically, the number of users in state j at time $t = 0$ is $K(\Pi_K(j, 0))$ and the multinomial probabilities are the j th row of $P(g, \Pi_K(0))$. Let $Y(M, \mu)$ be a multinomial random variable of size M with probability vector μ . Then the transition equation for the empirical distribution at time $t = 1$ for finite K can be written as

$$\Pi_K(1) \approx \frac{1}{K} \sum_{j=1}^{|X|} Y(K\Pi_K(j, 0), P(g, \Pi_K(0))), \quad (12)$$

which converges to the following equation as $K \rightarrow \infty$,

$$\pi(1) = \pi(0)P(g, \pi(0)). \quad (13)$$

Now that we have shown convergence of $\pi(1)$ it is straightforward to apply the same argument inductively to show that $\pi(t+1) = \pi(t)P(g, \pi(t))$ for all t .

Proposition 2: The equation $\pi(t+1) = \pi(t)P(g, \pi(t))$ has a unique fixed point and therefore the empirical distribution $\pi(t)$ has a unique steady state distribution.

Proof: Earlier in the paper (4), we defined a vector $h(g) \in \mathbb{R}^{|X|}$ such that the j th component of $h(g)$ defined the received power from a mobile in state j . If we examine the limiting SIR in (4) we see that the SIR for each user is determined by the power control g and the expected value of $h(g)$ computed with respect to $\pi(t)$ (i.e. the expected value is simply $\pi(t)h(g)$). The authors of [11] refer to this computation as the expectation of the effective interference from each user. Though we use the matched filter receiver as an example, this result extends to the MMSE receiver.

Hence if the SIR per user depends only on g and $\pi(t)h(g)$ we can write the equation in Proposition 1 as

$$\pi(t+1) = \pi(t)P(g, \pi(t)) \quad (14)$$

$$= \pi(t)P(g, \pi(t)h(g)). \quad (15)$$

Therefore a fixed point of (14) will also be a fixed point of (15).

Hence our fixed point equation of interest is $\pi = \pi P(g, \pi h(g))$. Recall that the vector $h(g)$ is the received power from each user and is bounded above and below by $h_{\min} \leq h(g) \leq h_{\max}$. Then for any constant c such that $h_{\min} \leq c \leq h_{\max}$ we can solve the linear equation

$$\pi(c) = \pi(c)P(c), \quad (16)$$

for $\pi(c)$. In our system model it can be shown that $P(c)$ is irreducible, and therefore the solution to (16) is unique. Define a function $\beta(c)$ such that

$$\beta(c) = \pi(c)h - c. \quad (17)$$

We can also show that on each end of the interval $[h_{\min}, h_{\max}]$ we have $\beta(h_{\min}) > 0$ and $\beta(h_{\max}) < 0$. Since $P(c)$ (and therefore $\beta(c)$) is a continuous function, we know there must be a zero of $\beta(c)$ on the interval $[h_{\min}, h_{\max}]$. Furthermore, for the formulation in this paper it can be shown that $\beta(c)$ is strictly decreasing. Therefore $\beta(c)$ has a unique zero at some point $c^* \in [h_{\min}, h_{\max}]$.

We have a unique c^* that implies $\beta(c^*) = \pi(c^*)h - c^* = 0$. Then by the irreducibility of the matrix $P(c)$ we know there exists a unique solution to the fixed point equation

$$\pi(c^*) = \pi(c^*)P(c^*) \quad (18)$$

$$= \pi(c^*)P(\pi(c^*)h), \quad (19)$$

where the second equality comes from $\beta(c^*) = 0$. But (19) is exactly what we were trying to show, hence there exists a unique solution to the fixed-point equation $\pi = \pi P(g, \pi h(g))$.

B. Value Function Simplification

Now that we have shown the existence of a unique fixed point for the empirical distribution of users in a CDMA cell, we can greatly simplify the value function (8).

Proposition 4: The value function proposed in equation (8) simplifies to the vector product $\pi(\infty, g)r(g)^T$, where $\pi(\infty, g)$ is the steady state value of the empirical distribution as a function of the control policy g .

Proof: Define a mapping $T^n(g, \pi)$ such that $T(g, \pi) = \pi P(g, \pi)$, $T^2(g, \pi) = (\pi P(g, \pi))P(g, \pi P(g, \pi))$, etc., which maps a starting vector $\pi(0)$ n steps into the future using (15). Then our value function as $K, N \rightarrow \infty$ is

$$V_\infty(g) = \lim_{T \rightarrow \infty} \lim_{K \rightarrow \infty} \left[\frac{1}{T+1} \sum_{n=0}^T E_{X, \pi} [R(g)] \right] \quad (20)$$

$$= \lim_{T \rightarrow \infty} \left[\frac{1}{T+1} \sum_{n=0}^T T^n(g, X(0)) r(g)^T \right] \quad (21)$$

$$= \lim_{T \rightarrow \infty} \left[\frac{1}{T+1} \sum_{n=0}^T T^n(g, X(0)) \right] r(g)^T \quad (22)$$

$$= \pi(\infty, g)r(g)^T, \quad (23)$$

where the last equality comes from the existence of a unique fixed point of the mapping T .

C. Optimization Formulation

Now that we have simplified our value function and dynamic equations we can formulate an optimization problem that minimizes our value function subject to constraints on system behavior. Recall that our original intent was to find the optimal power control policy subject to a constraint on the probability of packet loss due to deadline expiration. One way to do this is to constrain the frequency of visits to states where a deadline expiration occurs. We will represent this constraint (or possibly a set of constraints) through a matrix A and vector b such that our performance constraints are met if $\pi(\infty, g)A \leq b$. With the performance constraint and the development in the previous sections we can use the following non-linear program to solve for the power control policy that minimizes our value function (20):

$$\min_g \pi(\infty, g)r(g)^T \quad (24)$$

subject to:

$$\pi(\infty, g) = \pi(\infty, g)P(\pi(\infty, g))$$

$$\pi(\infty, g)\mathbf{1}^T = 1$$

$$\pi(\infty, g)A \leq b$$

$$\pi(\infty, g) \geq 0$$

Intuitively, we are minimizing the value function subject to the constraint that the steady-state empirical distribution $\pi(\infty, g)$ satisfies the fixed point equation (19). As well, $\pi(\infty, g)$ must be a probability distribution and satisfy our performance constraints. We should point out that this optimization problem is almost identical to that from [5], which considered a single user formulation. The only (and obviously critical) difference is the non-linear fixed point equation for the steady state of the system.

D. Technical Details

In the previous sections we did not discuss a few key technical questions. First, can we always find a solution to our proposed non-linear program? Second, were we justified in making this limiting approximation in the first place?

Unfortunately, it is possible to show that our proposed non-linear program is not convex. Clearly this is a significant drawback to our formulation since we are not guaranteed to find the optimal power control policy when we solve the optimization problem (24). However, one can show that it is possible to get arbitrarily close to the optimal solution by solving a sequence of linear programs. A detailed discussion of this proof is reserved for [7].

The answer to the second question requires two parts, addressing both the theoretical and practical aspects of our approximation. The critical theoretical question is whether or not our exchange of limits in (20) was valid. In our limiting value function $V_\infty(g)$ we first took the limit as $K \rightarrow \infty$ and then computed the value function using the dynamics for this limiting system. From a theoretical standpoint it would be reassuring if we could show that the value functions for the finite K systems, $V_K(g)$, converge to $V_\infty(g)$. Though we cannot compute the finite K value functions directly due to the complexity issues mentioned above, we can indeed show that $V_K(g) \rightarrow V_\infty(g)$. Once again, a detailed proof is reserved for the journal version [7] of this paper. The practical issue, which we address in the next section, is the similarity in performance of a finite user system to the theoretical optimum of the limiting system.

IV. NUMERICAL EXAMPLES

In this section we consider a numerical example of the optimization formulation (24) as well as some simulation data that checks the validity of our approximation. Of particular interest is the behavior of finite K systems operating under the control from the limiting system. Our results show that the approximation works well for a reasonable number of users.

A. System Model

We consider a specific instance of the model proposed in Section II. We have a group of users transmitting deadline sensitive data to a base station. Packets generated at each user announce a deadline $d = 100\text{ms}$ in the future, if a packet is not successfully transmitted by its deadline it is dropped. In each time slot a user may choose a transmitter power between .02 and .8 Watts in increments of .5dB. The wireless channel is modeled as a DTMC approximation to a log-normal shadowing channel. As suggested by Gudmundson [4], we model the log-normal shadowing as a Gaussian process in dB units. In order to fully describe the shadowing model we must specify the mobile velocity, mean path loss, standard deviation, and a correlation coefficient. Once these parameters are determined we split the range of channel gains into a finite number of bins and construct a DTMC that approximates the Gaussian process using the method suggested in [10]. The following table contains the channel gain regions for the DTMC and other channel parameters.

Path Loss Ranges in dB	$(-\infty, 112.5), (112.5, 117.5), (117.5, 122.5), (122.5, 127.5), (127.5, \infty)$
Mobile Velocity	50km/h
Mean Path Loss	120dB
Shadowing Std. Dev	$\sigma = 4.3\text{dB}$
Shadowing Correlation	$\zeta_D = .3$ for $D = 20\text{m}$
Noise Threshold N_0	-150dBW
Carrier Frequency	900MHz

We assume the linear receiver structure follows the matched filter proposed in Section II. We model the probability of a user

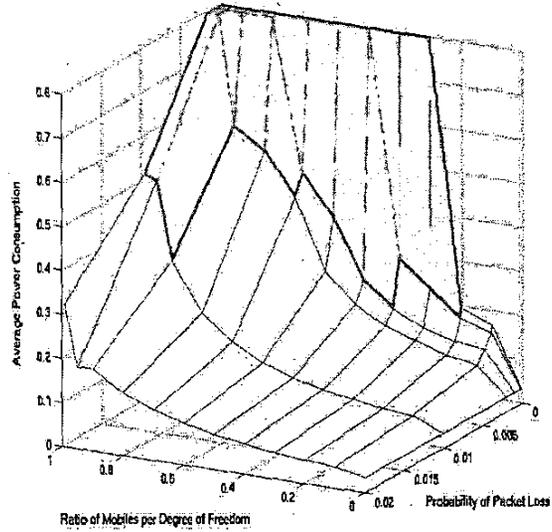


Fig. 1. Optimal Average Power vs. α and Probability of Deadline Expiration. The problem is infeasible in the shaded region.

failing to transmit a packet as a function of that user's SIR [13],

$$P(x_i(t+1) = x_i(t) | SIR_i(g, t)) = 1 - \exp^{-\frac{F_i}{\alpha}}, \quad (25)$$

where F_i represents the fade margin [3] and can be computed as $F_i = \beta F_0 SIR_i$, where F_0 is a constant and β denotes the efficiency of the power control ($\beta = 1$ corresponds to perfect control). In this example we set $F_0 = 7.5$ and $\beta = 1$.

Given this setup our goal is to minimize average transmission power subject to a constraint on the fraction of packets lost to deadline expiration. We will consider constraints on the probability of packet loss ranging from .005% to 2% and values of α ranging from 0 to 1. The constraint values correspond to relatively tight to quite loose constraints on packet loss. Recall the parameter α is the ratio of the number of users in the system to the length of the spreading sequences. Therefore values of α close to zero correspond to a lightly loaded system with the load increasing substantially as α approaches 1.

B. Optimization Results

First we consider the solution of the non-linear program (24) for various values of α and our constraint on probability of deadline expiration. Figure 1 plots the average transmission power vs. α and probability of deadline expiration. The shaded area denotes the infeasible region of the non-linear program. Within the infeasible region the optimal control that minimized the infeasibility in the constraints is plotted.

The optimal average power in Figure 1 is fairly intuitive. For relatively loose constraints on probability of packet loss we are willing to tolerate more interference. This means we can afford to transmit with higher power as the numbers of users per degree of freedom (i.e. α) increases. However, as the performance constraints become tighter we become less tolerant of interference and cannot afford to raise our transmission power as α increases. Therefore a system with tight deadline constraints cannot support as many users as a lossy system, even if the users in the tight system can transmit with fairly high power. It is also interesting to note that the control policy that minimizes the infeasibilities in the shaded region is to transmit with maximum power.

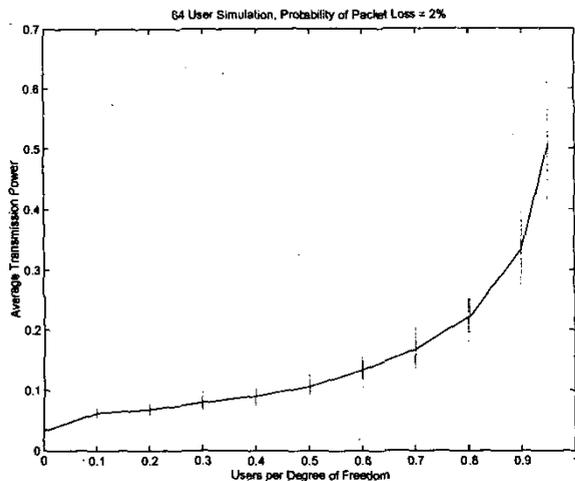


Fig. 2. Simulation Results for a 64 User System

C. Simulation Results

Now we consider the behavior of a finite number of users following the optimal control for the limiting system. The plot in Figure 2 shows data taken from simulation runs of 64 users following the optimal limiting control for 2% probability of deadline expiration. The solid line plots the mean value taken from Figure 1, and the points around the line plot 100 samples of the mean transmission power that were taken after the simulation reached steady state. The standard deviation from the limiting value is between 5-10% of the mean, and does not change significantly for different values of α . So, for this example, our approximation performs quite well.

V. CONCLUSIONS

We have developed a general framework for optimizing link adaptation in CDMA systems in the wideband limit. We then applied this framework to determine optimal power control to minimize transmit power under a data delay constraint. The crux of this problem is accounting for the responsive interference in CDMA systems. In general, the optimal adaptation policy is difficult to find since the size of the problem grows exponentially with the number of users in the system. However, we showed that in the limiting regime of long spreading sequences and large numbers of users, the size of the optimal adaptation problem collapses to that of a single user problem, but with the added complication of non-linear dynamics. This substantial reduction in complexity is critical for solving useful multi-user control problems, which previously required a substantial reduction in modeling detail in order to make the problem tractable. Finally, we presented simulation results that showed the optimal power control from the limiting regime provides a reasonable approximation for optimal power control in a finite user system.

This formulation can also be used with much more general link adaptation controls and system models. In previous publications we considered single user formulations for optimal joint source-channel coding, systems with delayed and inaccurate channel estimates, as well as the effect of mobile speed and propagation environment on optimal link controls. All of these results (and many others) can be extended to multi-user CDMA systems using the methods proposed in this paper.

REFERENCES

[1] N. Bambos, S. Kandukuri, *Power Controlled Multiple Access in Wireless Communications Networks*, IEEE INFOCOM 2000.

[2] S. Gitisen, N. Bambos, *Power-Controlled Data Prefetching/Caching in Wireless Packet Networks*, IEEE INFOCOM 2002.
 [3] V. Garg and E. Sneed, *Digital Wireless Local Loop Systems*, IEEE Communications Magazine, pp. 112-115, Oct. 1996.
 [4] M. Gunmundson, *Correlation Model for Shadow Fading in Mobile Radio Systems*, Electronics Letters, Vol. 27, November 1991.
 [5] T. Holliday, A. Goldsmith, and P. Glynn, *Optimal Power Control and Source-Channel Coding for Delay Constrained Traffic over Wireless Channels*, IEEE ICC 2002.
 [6] T. Holliday, A. Goldsmith, and P. Glynn, *Wireless Link Adaptation Policies: QoS for Deadline Constrained Traffic with Imperfect Channel Estimates*, IEEE ICC 2002.
 [7] T. Holliday, A. Goldsmith, and P. Glynn, *Asymptotically Optimal Link Adaptation for WCDMA Systems*, Journal version of Globecom2002 submission.
 [8] Y.Y. Kim, S. Li, *Modeling Fast Fading Channel Dynamics for Packet Data Performance Analysis*, Proceedings IEEE INFOCOM '98.
 [9] U. Madhow and M. Honig, *MMSE Interference Suppression for Direct-Sequence Spread-Spectrum CDMA*, IEEE Transactions on Communication, vol. 42, pp.3178-3188, Dec. 1994.
 [10] G. Stuber, *Principles of Mobile Communication*, Kluwer Academic Publishers, 1999.
 [11] D. Tse and S. Hanly, *Linear Multiuser Receivers: Effective Interference, Effective Bandwidth and User Capacity*, IEEE Transactions on Information Theory, v.45, No. 2, Mar. 1999, pp. 641-657.
 [12] A. Ward and P. Glynn, *Weak Convergence of the Steady State Distribution for a Sequence of Regenerative Processes*, Submitted for publication.
 [13] M. Zorzi, R. Rao, and L. Milstein, *Error Statistics in Data Transmission over Fading Channels*, IEEE Transactions on Communication, vol. 46, pp. 1468-1477, Nov 1998.