Distributed power control and spreading gain allocation in CDMA data networks*

Seong-Jun Oh
EECS Department
University of Michigan
Ann Arbor, MI 48109

Tava Lennon Olsen
IOE Department
University of Michigan
Ann Arbor, MI 48109

Kimberly M. Wasserman†
EECS Department
University of Michigan
Ann Arbor, MI 48109

Abstract

We study the radio resource allocation problem of distributed joint transmission power control and spreading gain allocation in a DS-CDMA mobile data network. The network consists of K base stations and M wireless data users. The data streams generated by the users are treated as best-effort traffic, in the sense that there are no prespecified constraints on the quality of the radio channels. We are interested in designing a distributed algorithm that achieves maximal (or near-maximal in some reasonable sense) aggregate throughput, subject to peak power constraints. We provide an algorithm where neighboring base stations coordinate in a distributed fashion to control the powers and spreading gains of the users, and show that it converges to a Nash equilibrium. In general, there may be multiple equilibria; however, certain structural properties of the throughput expression can be exploited to significantly trim the search space. The numerical results indicate that with these modifications, the algorithm frequently converges in just a few iterations to the throughput maximizing (globally optimal) power and spreading gain allocation.

1 Introduction

The tremendous success of voice-based cellular telephone networks has spawned an increasing interest in mobile wireless data communication. Indeed, a number of wireless data networks have recently emerged in the marketplace spanning the domains of cellular networks (e.g., CDPD, GPRS, and EDGE) to metropolitan, local area, and ad hoc networks operating in the ISM band[1] (e.g., ARDIS, Metricom Richochet, Cisco Aironet, AT&T WaveLAN, Utilicom Longranger, Rooftop, Bluetooth, and HomeRF). As compared to voice traffic, data traffic is less senstive to delays, but more sensitive to transmission errors; reliability is assured via retransmissions. From a network layer perspective, data communication is often treated as a best-effort service, and therefore achieving high aggregate data throughput is one of the primary goals. Towards this, the delay tolerance of data traffic can

*A portion of this work was presented at the IEEE Conf. Computer Commun. (INFOCOM'00), Tel Aviv, Israel, 2000. This work was supported in part by the U.S. Army Research Office under grant DAAH04-96-1-0177.

†Corresponding author: K. M. Wasserman, EECS Department, University of Michigan, 1301 Beal Avenue, Ann Arbor, MI 48109-2122. (e-mail: wass@eecs.umich.edu).

[1]Industrial, Scientific, and Medical band (902–928 MHz, 2400–2483.5 MHz, and 5725–5850 MHz) approved by the FCC for unlicensed use under certain restrictions.

be exploited to allow the design of spectrum-efficient radio resource allocation algorithms adapting the received energy per bit (ratio of the received power per bit to the transmission rate) to the current interference and channel conditions.

In this paper, we study the radio resource allocation problem of distributed power control and spreading gain (transmission rate) allocation in a Direct-Sequence Code Division Multiple Access (DS-CDMA) mobile data network. The network consists of M wireless data users and K stationary base stations (access points). At any given point in time, each user is connected, that is, assigned, to at most one base station; however, each base station is capable of simultaneously receiving transmissions from several different users. The users generate non-real-time streams of data packets (e.g., paging, electronic mail, facsimile, file transfer, etc.), and transmit to appropriately chosen base stations using different spreading codes over a common wideband radio channel. We are interested in designing a joint transmission power control and spreading gain allocation algorithm for the users that achieves maximal (or near maximal in some reasonable sense) aggregate throughput, subject to peak power constraints. We focus on the uplink (user-to-base station), but the downlink can be treated similarly.

In terms of maximizing aggregate throughput, or any other global network-wide objective, it is best for the base stations fully coordinate in order to determine the optimal allocations; but this is, of course, not practical. In [8], a radio resource allocation algorithm is provided for a DS-CDMA network in which the base stations act autonomously, that is, the base stations do not coordinate (except for soft handoffs, etc.), to control the transmission powers and spreading gains of the users. Each base station attempts to maximize the aggregate throughput of just the users in its cell, subject to constraints on peak power and, to prevent excessive intercell interference, total received power. The simulations of the algorithm presented in [8] show that it outperforms other autonomous resource allocation schemes, such as equal received power and round robin, with respect to throughput and average delay, even for users with low quality channels.

With a moderate increase in complexity, even higher performance gains would be expected by allowing a limited amount of coordination between base stations. The basic idea is that light loading or low interference in neighboring cells can be exploited via the exchange of information between neighboring base stations to achieve a higher network spectral efficiency than can be achieved via less complex greedy localized maximizations by each base station separately. In this paper, we are interested in developing an algorithm where neighboring base stations coordinate in a distributed fashion to jointly control the transmission powers and spreading gain allocations of the users.

The problem of distributed power control in wireless communication networks has received considerable attention in the past. For a network in which the QoS requirement of each user is a prespecified target $\mathcal{E}_b/I_0$ value, Hanly [1] and Yates and Huang [2] have developed distributed power control and base station assignment algorithms based on interference and channel measurements local to the base stations, that minimize the total transmitted power subject to target $\mathcal{E}_b/I_0$ values, and provided conditions for their convergence. Grandhi et al. [3] have developed a distributed power control algorithm that maximizes the minimum $\mathcal{E}_b/I_0$ value of the users, subject to constraints on peak transmission power. For a network in which each user generates a bursty traffic stream, and specifies its QoS requirement in terms of the probability that the $\mathcal{E}_b/I_0$ falls below a target level, Mitra and Morrison [4] have developed a distributed power control algorithm based on measurements of the mean and variance of the interference at the base stations, and provided conditions for its convergence. For a network in which each user generates a variable rate traffic stream, and specifies its QoS requirement in terms of a target $\mathcal{E}_b/I_0$ value, Kim [5] has developed a distributed rate-regulated power control algorithm. Ji and Huang [6] and Feng et al. [7] have studied the power control problem from a game-theoretic perspective, and developed a power control algorithm, which under certain assumptions, converges to a Nash equilibrium.

Here, data communication is treated as a best-effort service, in the sense that there are no prespecified target $\mathcal{E}_b/I_0$ values for a user; reliability is assured via automatic repeat request (ARQ) protocols. As stated above, our objective is to develop a distributed joint power control and spreading gain allocation algorithm that achieves high aggregate network throughput, subject to peak power constraints. We provide such an algorithm and show that it converges to a Nash equilibrium. In general, there may be multiple equilibria; however, the numerical results indicate that the algorithm frequently converges in just a few iterations to the throughput maximizing (globally optimal) power and spreading gain allocation. In addition, while the maximization of aggregate throughput is a computationally complex nonlinear programming problem, and, in fact, the only method guaranteed to determine the global optimum is exhaustive search, the algorithm has far less complexity. This has been achieved by exploiting certain structural properites of the throughput expression along with a heuristic derived from autonomous optimization to significantly trim the search space.

The remainder of the paper is organized as follows. In Section 2, we describe the system model. In Section 3, we formulate the optimization problem. In Section 4, we provide the distributed joint transmission power and spreading gain allocation algorithm. In Section 5, we present numerical results, and in Section 6, we conclude.

A word on notation used in this paper: Vector quantities are denoted by boldface letters.

2 System model

Define the channel gain vector $\boldsymbol{g}_i = (g_{i1}, g_{i2}, \ldots, g_{iK})$, where $g_{ik} \geq 0$ is the instantaneous gain of the radio channel between the i^{th} user and the k^{th} base station, $i = 1, 2, \ldots, M$, $k = 1, 2, \ldots, K$. When the channel gain between a user and a base station is less than some predefined $\epsilon > 0$, we take it to be zero. Define $a_i \in \{1, 2, \ldots, K\}$ to be the base station to which the i^{th} user is assigned, $i = 1, 2, \ldots, M$, and $\mathcal{C}_k = \{i : a_i = k\}$ the set of users assigned to the k^{th} base station, that is, in the k^{th} "cell", $k = 1, 2, \ldots, K$. Let $|\mathcal{C}_k|$ denote the cardinality of $\mathcal{C}_k$, $k = 1, 2, \ldots, K$. Without loss of generality, assume that if $i, j \in \mathcal{C}_k$, then $g_{ik} > g_{jk} > 0$, for $i < j$, so that among the users of the k^{th} cell, the lower the index the higher the channel quality, $k = 1, 2, \ldots, K$. Define $0 \leq p_i \leq p_{max}$ and $N_i \in \mathbb{R}_+$ to be the transmission power and spreading gain[2], respectively, of the i^{th} user, $i = 1, 2, \ldots, M$. Let $\boldsymbol{p}^k$ be the vector of powers of the users of the k^{th} cell, $k = 1, 2, \ldots, K$. Let b_i be the index of user i in $\mathcal{C}_{a_i}$ so that user i transmits with power $p_{b_i}^{a_i}$, $i = 1, 2, \ldots, M$. The receivers at the base station are of the conventional matched filter type.

The $\mathcal{E}_b/I_0$ experienced by the i^{th} user is equal to the product of the spreading gain and the SINR, and is given by

$$\Gamma_i = \frac{p_i g_{ia_i} N_i}{\alpha \sum_{j=1, j \neq i}^{M} p_j g_{ja_i} + \sigma_{a_i}^2}, \tag{1}$$

where $\sigma_{a_i}^2$ is the background interference and thermal noise power at the a_i^{th} base station, and

[2]It is difficult to implement a non-integer spreading gain; however, for analytical simplicity we assume that spreading gains can be real numbers. In practice, the spreading gain would be rounded to, say, the next highest integer.

$\alpha > 0$ is a constant[3] depending on the statistical characteristics between the spreading codes of users, $i = 1, 2, \ldots, M$. The probability $f : \mathbb{R}_+ \to [0, 1]$ of a successful packet transmission is a differentiable and non-decreasing function of the $\mathcal{E}_b/I_0$ that depends on factors such as the modulation/demodulation technique, interleaving depth, and the channel coding scheme.

3 Structural properties of optimal solutions

In this section, we summarize some earlier work on power control and variable spreading gain for DS-CDMA networks in [8], which will be of use in designing distributed resource allocation algorithms.

The instantaneous aggregate throughput in packets/sec/Hz/cell (channel efficiency) $S : \mathbb{R}_+^M \times \mathbb{R}_+^M \to \mathbb{R}_+$ is given by

$$\begin{aligned} S(\boldsymbol{p}, \boldsymbol{N}) &= \frac{1}{K} \sum_{i=1}^{M} \frac{\beta}{N_i} f(\Gamma_i) \\ &= \frac{1}{K} \sum_{i=1}^{M} \frac{\beta}{N_i} f\left(\frac{p_i g_{i a_i} N_i}{\alpha \sum_{j=1, j \neq i}^{M} p_j g_{j a_i} + \sigma_{a_i}^2} \right) \\ &= \frac{1}{K} \sum_{k=1}^{K} \sum_{i \in \mathcal{C}_k} \frac{\beta}{N_i} f\left(\frac{p_i g_{ik} N_i}{\alpha \sum_{j=1, j \neq i}^{M} p_j g_{jk} + \sigma_k^2} \right), \end{aligned} \tag{2}$$

where β is the ratio of the channel code rate used by the users to the number of bits per packet.

The jointly optimal power and spreading gain allocation is that which maximizes aggregate data throughput, and is obtained by solving the following optimization problem:

$$(P1) \qquad \max_{(\boldsymbol{p}, \boldsymbol{N}) \in \mathcal{F}} S(\boldsymbol{p}, \boldsymbol{N}),$$

where the feasible set $\mathcal{F}$ is given by

$$\mathcal{F} = \{(\boldsymbol{p}, \boldsymbol{N}) : 0 \leq p_i \leq p_{max} \text{ and } N_i \in \mathbb{R}_+, i = 1, 2, \ldots, M\}. \tag{3}$$

Note that the objective function in $(P1)$ is non-convex and the size of the feasible set explodes with M; solving for the globally optimal power and spreading gain allocation is computationally complex, and the only method guaranteed to find the global optimum is exhaustive search.

Since (2) is separable in $\boldsymbol{N}$, we simplify $(P1)$ by obtaining the optimal spreading gain allocations, which are determined by the channel gains and transmission powers of the users. The spreading gain of a user is the ratio of the signal bandwidth after spreading (i.e., the channel bandwidth or chip rate) to the bandwidth before spreading, or equivalently, the ratio of the time duration of a bit to the time duration of a chip (chips per bit). The spreading gain of a user can be controlled by dynamically increasing or decreasing the time duration of a transmitted bit (decreasing or increasing the transmission rate of the user), while maintaining a constant chip rate.

The optimal spreading gain control mechanism is given in the proposition below.

[3]Typical values are $\alpha = 1$ (synchronous) or $1/3$ (asynchronous).

Proposition 1 (Proposition 1 of [8]) The optimal spreading gain N_i^* of the i^{th} user is given by

$$N_i^* = \frac{\gamma^*(\sigma_{a_i}^2 + \alpha \sum_{j=1, j\neq i}^{M} p_j g_{j a_i})}{p_i g_{i a_i}}, \tag{4}$$

$i = 1, 2, \cdots, M$, where

$$\gamma^* = \arg\max_{\gamma \geq 1} \left\{ \frac{1}{\gamma} f(\gamma) \right\}. \tag{5}$$

For $\gamma \geq 1$ and the number of bits per packet large enough,[4] $f(\gamma)/\gamma$ is generally pseudo-concave. Thus the Karush-Kuhn-Tucker conditions are sufficient as well as necessary for optimality.

Thus the optimal spreading gain of a user is inverse linear in its SINR, or equivalently, since the spreading gain is inversely proportional to the transmission rate, the optimal transmission rate of a user increases linearly as its SINR increases.

Under the optimal spreading gain allocation in Proposition 1, the aggregate throughput (2) becomes

$$\begin{aligned} S(\boldsymbol{p}, \boldsymbol{N}^*) &= \frac{1}{K} \frac{\beta f(\gamma^*)}{\gamma^*} \sum_{i=1}^{M} \frac{p_i g_{i a_i}}{\alpha \sum_{j=1, j\neq i}^{M} p_j g_{j a_i} + \sigma_{a_i}^2} \\ &= \frac{1}{K} \frac{\beta f(\gamma^*)}{\gamma^*} \sum_{k=1}^{K} \sum_{i \in \mathcal{C}_k} \frac{p_i g_{ik}}{\alpha \sum_{j=1, j\neq i}^{M} p_j g_{jk} + \sigma_k^2}. \end{aligned} \tag{6}$$

Thus the optimization problem $(P1)$ reduces to one of determining an optimal transmission power allocation. Note that maximizing (6) is equivalent to maximizing the sum of the SINRs. In addition, since the transmission rates are given by the ratio of the spreading bandwidth (which is constant) to the spreading gains, and the optimal spreading gains are given by the ratio of the constant γ^* to the SINRs, maximizing (6) is also equivalent to maximizing the sum of the transmission rates.

The optimization problem $(P1)$ under the optimal spreading gain allocation (4) becomes

$$(P2) \qquad \max_{\boldsymbol{p} \in \mathcal{P}} S(\boldsymbol{p}, \boldsymbol{N}^*),$$

where the feasible set $\mathcal{P}$ is given by

$$\mathcal{P} = \{\boldsymbol{p} : 0 \leq p_i \leq p_{max},\ i = 1, 2, \ldots, M\}. \tag{7}$$

Note that although (6) is convex in each p_i, it is not convex in $\boldsymbol{p}$, and thus the Karush-Kuhn-Tucker conditions are not sufficient for optimality (i.e., we cannot guarantee global optimality of a solution obtained by taking partial derivatives). Thus we must optimize a non-convex function over a space of size 2^M. However, as the following two lemmas show, the feasible set $\mathcal{P}$ can be significantly trimmed without loss of optimality. The first lemma shows that $\mathcal{P}$ can be reduced to the finite set of its vertices.

Lemma 1 (Lemma 1 of [8]) If $\boldsymbol{p} = (p_1, p_2, \cdots, p_M)$ solves $(P2)$, then $p_i = 0$ or $p_i = p_{max}$ for all $i \in \{1, 2, \ldots, M\}$.

[4]We assume that the number of bits per packet is large enough so that $\gamma \geq 1$ in (5) avoids the degeneracy of $f(\gamma)/\gamma$ tending to infinity as γ tends to zero, which arises from the fact that the bit error probability does not tend to zero as the bit energy-to-interference density ratio tends to zero.

In Lemma 2 below, we show that under the optimal power allocation, a given user in an arbitrary cell transmits only if all users with higher channel gains in that cell transmit; equivalently, if a given user is not allowed to transmit (transmission power zero) then neither is any other user in the same cell with lower channel gain. This lemma can be deduced from Lemma 2 of [8] and Lemma 1 above.

Lemma 2 If $\boldsymbol{p} = (p_1, p_2, \cdots, p_M)$ solves $(P2)$, then since $g_{ik} > g_{jk}$, $i < j$, $i, j \in \mathcal{C}_k$, if $p_j = p_{max}$ for some $j \in \mathcal{C}_k$, then $p_i = p_{max}$ for all $i < j$, $i, j \in \mathcal{C}_k$, $k = 1, 2, \ldots, K$.

Define $\boldsymbol{\pi}^{(k,i)}$ to be a vector of dimension $\mathcal{C}_k$ with the first i entries being equal to p_{max} and the last $|\mathcal{C}_k| - i$ entries being equal to zero, $i = 0, 1, \ldots \mathcal{C}_k$, $k = 1, 2, \ldots, K$. Define $\diamond^k = \{\boldsymbol{\pi}^{(k,i)}; i = 0, 1, \ldots |\mathcal{C}_k|\}$, $k = 1, 2, \ldots, K$. Proposition 1 and Lemmas 1 and 2 are used to design a simple hybrid CDMA/TDMA resource allocation strategy solving $(P1)$, subject to peak power constraints as follows. First define γ^* by (5). Next, set the powers of the users in the k^{th} cell equal to $\boldsymbol{p}^k$, $k = 1, 2, \ldots, K$, which are found by solving

$$(P3) \qquad \max_{\substack{\boldsymbol{p}^k \in \{0, p_max\}^k \\ k=1,2,\ldots,K}} \frac{1}{K} \frac{\beta f(\gamma^*)}{\gamma^*} \sum_{k=1}^{K} \sum_{i \in \mathcal{C}_k} \frac{p_{b_i}^k g_{ik}}{\alpha \sum_{j=1, j\neq i}^{M} p_{b_j}^{a_j} g_{jk} + \sigma_k^2}.$$

Last, set the spreading gain of the i^{th} user to

$$N_i^* = \frac{\gamma^* (\sigma_{a_i}^2 + \alpha \sum_{j=1, j\neq i}^{M} p_{b_j}^{a_j} g_{j a_i})}{p_{b_i}^{a_i} g_{i a_i}},$$

$i = 1, 2, \cdots, M$.

Note that the users can compute their optimal spreading gain allocations in a simple decentralized fashion using information local to the base stations: Each base station measures the SINRs and transmits them to the users in their cells using pilot tones and control channels; the users invert their respective SINRs and multiply by the constant γ^* as in (4). To decentralize $(P1)$, it then remains to decentralize the optimization problem $(P3)$, that is, to design a distributed transmission power allocation algorithm that solves (or nearly solves in some reasonable sense) $(P3)$.

4 Distributed power control algorithm

In this section, we provide a distributed implementation of the optimization in $(P3)$. Let $\mathcal{I}_k = \{j \neq k : g_{ik} > 0 \text{ for some } i \in \mathcal{C}_k\}$ be the neighbors of the k^{th} cell, that is, the set of cells in which the users of the k^{th} cell may cause interference, $k = 1, 2, \ldots, K$. Depending on the cell size and p_{max}, $\mathcal{I}_k$ typically consists of only the cells adjacent to the k^{th} cell, $k = 1, 2, \ldots, K$. Let $I_k(\boldsymbol{p}) = \sum_{i=1}^{M} g_{ik} p_i + \sigma_k^2$ be the total interference and noise power at the k^{th} base station, $k = 1, 2, \ldots, K$.

Let $J(\boldsymbol{p}) = S(\boldsymbol{p}, \boldsymbol{N}^*)$. In the following expression, let $\boldsymbol{y}^k$ denote a "dummy" vector of powers of the users in the k^{th} cell, and $\boldsymbol{p}^{-k} = (\boldsymbol{p}^1, \boldsymbol{p}^2, \ldots, \boldsymbol{p}^{k-1}, \boldsymbol{p}^{k+1}, \ldots, \boldsymbol{p}^K)$ the vector of all power vectors except the that of the k^{th} cell. Let $(\boldsymbol{p}^{-k}, \boldsymbol{y}^k) = (\boldsymbol{p}^1, \boldsymbol{p}^2, \ldots, \boldsymbol{p}^{k-1}, \boldsymbol{y}^k, \boldsymbol{p}^{k+1}, \ldots, \boldsymbol{p}^K)$. Define the function $J_k : [0, p_{max}]^M \to \mathbb{R}_+$ as

$$J_k((\boldsymbol{p}^{-k}, \boldsymbol{y}^k)) = \sum_{i \in \mathcal{C}_k} \frac{g_{ik} y_{b_i}^k}{I_k((\boldsymbol{p}^{-k}, \boldsymbol{y}^k)) - g_{ik} y_{b_i}^k} + \sum_{j \in \mathcal{I}_k} \sum_{m \in \mathcal{C}_j} \frac{g_{mj} p_{b_m}^j}{I_j((\boldsymbol{p}^{-k}, \boldsymbol{y}^k)) - g_{mj} p_{b_m}^j}, \qquad (8)$$

$k = 1, 2, \dots, K$. The first term in (8) is the throughput attained in the k^{th} cell and accounts for intracell interference; the second term is the throughput attained in the neighboring cells of the k^{th} cell, and since the transmission powers of the users in the k^{th} cell only appear in the denominator of the summands, accounts for intercell interference. Note that J_k is a partial sum of J, and that J_k and J_ℓ may share some common terms, $k \neq \ell$, $k, \ell = 1, 2, \dots, K$. From an M-person game-theoretic perspective, $\{J_k\}_{k=1}^K$ may be viewed as a set of utility functions. These utility functions, however, do not satisfy the assumptions required in [6] or [7]; in particular, J_k is neither decreasing nor quasi-concave in $\boldsymbol{p}^k$, $k = 1, 2, \dots, K$.

Now consider the following distributed power allocation algorithm. Starting from an arbitrary but feasible initial state, and proceeding in a round robin fashion among base stations, the powers at the k^{th} step of the n^{th} iteration are given by

$$\boldsymbol{p}^k(n) = \arg \max_{\boldsymbol{y}^k \in \{0, p_{max}\}^{|\mathcal{C}_k|}} \left\{ J_k((\boldsymbol{p}(n-1)^{-k}, \boldsymbol{y}^k)) \right\} \tag{9}$$

$$\boldsymbol{p}^\ell(n) = \boldsymbol{p}^\ell(n-1), \tag{10}$$

$n \in \mathbb{Z}_+$, $k \neq \ell$, $k, \ell = 1, 2, \dots, K$. During the k^{th} step, only the powers of the users in the k^{th} cell may be adjusted. At the end of the k^{th} step of each iteration, the k^{th} base station notifies the $(k+1)^{st}$ base station, which then polls its neighboring cells to obtain the necessary information (discussed below) to begin the $(k+1)^{st}$ step.

Note that the structural properties of the globally optimal power allocation provided in Lemmas 1 and 2 were exploited to significantly trim the search space and induce an ordering on the users of each cell, which greatly simplify computation of (9) and thus implementation of the algorithm. First, by Lemma 1, the base stations do not need to consider the full range of transmission powers for each user in their cells when computing (9), but just the extreme values (zero and p_{max}). Second, by Lemma 2 it is optimal for each base station to allocate powers to the users in its cell in decreasing order of their channel gains. Thus the k^{th} base station does not need to consider all $2^{|\mathcal{C}_k|}$ power allocations, but only $|\mathcal{C}_k| + 1$ power allocations, namely, those in Π^k, $k = 1, 2, \dots, K$.

Note that no base station has control over the transmission power of the users in any cell other than its own; however, each base station appropriately balances a user's desire for a high transmission power with the amount of interference it will generate to other users. Furthermore, replacing J_k with J in (8) results in the same sequence of transmission powers, but requires global information regarding all the channel gains and interference levels. This replacement is possible because J_k includes all the information regarding the impact of and to the users in the k^{th} cell on and by the rest of the network, $k = 1, 2, \dots, K$. Also note that parallelism is possible among cells who are not mutually interfering. Specifically, if $\mathcal{I}_k \cap \mathcal{I}_\ell = \emptyset$, then the k^{th} and ℓ^{th} cells may adjust the transmission powers of their respective users simultaneously, $k, \ell = 1, 2, \dots, K$. Practically, the cells can be divided into clusters, and appropriate cells in disjoint clusters may execute the algorithm simultaneously, where the size of cluster is given by the frequency reuse factor, typically four or seven, of a time- or frequency-slotted system.

Definition 1 A transmission power allocation $\boldsymbol{p}$ is a Nash equilibrium point if it satisfies

$$\boldsymbol{p}^k = \arg \max_{\boldsymbol{y}^k \in [0, p_{max}]^{|\mathcal{C}_k|}} J_k((\boldsymbol{p}^{-k}, \boldsymbol{y}^k)), \tag{11}$$

$k = 1, 2, \dots, K$.

A Nash equilibrium is therefore a power vector such that no base station can improve its utility by changing just the powers of the users in its cell.

Proposition 2 The distributed power allocation algorithm converges to transmission power allocation that is a Nash equilibrium.

Proof Since J_k can be replaced by J in (9), $k = 1, 2, \ldots, K$, the sequence of total cost $\{J(\boldsymbol{p}(n)); n \in \mathbb{Z}_+\}$ is nondecreasing in n. Furthermore, $\{J(\boldsymbol{p}(n)); n \in \mathbb{Z}_+\}$ is bounded above. Therefore it converges to $J(\hat{\boldsymbol{p}}) < \infty$ for some $\hat{\boldsymbol{p}} \in [0, p_{max}]^M$ such that $\hat{\boldsymbol{p}}^k \in \Pi^k$, $k = 1, 2, \ldots, K$. Again, since J_k can be replaced by J in (9), $k = 1, 2, \ldots, K$, $\hat{\boldsymbol{p}}$ is a Nash equilibrium point. ■

Once the power control algorithm terminates, the base stations transmit the resulting SINRs to the users in their cells so they can set their spreading gains according to (4). This assumes the algorithm operates at a faster time scale than the coherence time of the most rapidly varying channel, so that the gains can be accurately tracked and are effectively constant during the algorithm execution time. We may alternatively assume that the channel gain measurements are averaged over the short-term (i.e., Rayleigh) multipath fading, so that the algorithm tracks the long-term fading (i.e., log-normal shadowing), which is assumed constant over the algorithm execution time.

One question that arises is the following: How distributed is this algorithm? The algorithm requires coordination between neighboring base stations, which is more complex than an autonomous implementation but considerably less complex than a centralized implementation; indeed, in a large network, $|\mathcal{I}_k| \ll K$, $k = 1, 2, \ldots, K$. However, if the neighboring cells, or more precisely, the interference levels caused to the users in these cells, are not explicitly taken into account, then there is no incentive for the base station not to maximize the aggregate throughput in its cell. Under the distributed algorithm, the k^{th} base station requires knowledge of the same information as required by the algorithm in [1], namely, g_{ik}, $g_{i\ell}$, $I_k(\boldsymbol{p})$, $I_\ell(\boldsymbol{p})$, $i \in \mathcal{C}_k$, $\ell \in \mathcal{I}_k$, plus knowledge of $g_{j\ell}$ and $\boldsymbol{p}^\ell$, $j \in \mathcal{C}_\ell$, $\ell \in \mathcal{I}_k$ $k = 1, 2, \ldots, K$. This information can be obtained via control channels and pilot tones. Similar quantities of information are also considered in [9].

Two other questions are: How close does the distributed power allocation algorithm come to the solution of $(P3)$? And what is its rate of convergence? Since there are no prespecified $\mathcal{E}_b/I_0$ constraints, the algorithm always converges to feasible power allocation. In addition, the performance attained by the intermediate power allocations increases with each step of each iteration (see proof of Proposition 2); therefore, the longer the algorithm is allowed to execute, the higher the resulting throughput. If the Nash equilibrium is unique, then since J_k can be replaced by J, $k = 1, 2, \ldots, K$, the equilibrium is not only Pareto optimal[5], but is globally optimal in the sense of matching the solution to $(P3)$ (and thus $(P2)$). However, the determinant of the second principle minor of the Hessian formed by J is easily shown to be negative, and thus $J(\boldsymbol{p})$ is neither concave nor convex in $\boldsymbol{p}$. In fact, there may exist multiple Nash equilibria (see below) in which case the distributed power allocation algorithm is not a contraction mapping, and $\{J_k\}_{k=1}^K$ are not quasi-concave. These two facts make it difficult to determine how close the distributed algorithm comes to the optimal solution and its rate of convergence. However, the numerical results indicate that by exploiting Lemmas 1 and 2 to trim the search space and induce an ordering on the users of each cell as described earlier, the algorithm frequently converges in just a few iterations to the

[5]A power vector is Pareto optimal if for any other power vector, if one base station improves its utility, then there is at least one other base station whose utility decreases.

throughput maximizing (globally optimal) power and spreading gain allocation.

Remark 1 [Multiple Nash equilibria] In general, there may be several transmission power allocations that are Nash equilibria. For example, consider a network consisting of two users ($M = 2$) and two base stations ($K = 2$). Let $p_{max} = 10$ mW, $g_{11} = g_{12} = 8 \times 10^{-4}$, $g_{22} = g_{21} = 6 \times 10^{-4}$, and $I_1 = I_2 = 10^{-4}$. It can easily be seen that $\boldsymbol{p} = (10, 0)$ is a Nash equilibrium with $J(\boldsymbol{p}) = 80$. However, $\boldsymbol{p}' = (0, 10)$ is also a Nash equilibrium with $J(\boldsymbol{p}') = 60$.

5 Numerical results

In this section, we present numerical results from computer simulations evaluating the performance (with respect to ($P3$)) and convergence rate of the two distributed transmission power algorithms. We also compare our results with the autonomous algorithm from [8]. Four different cases are used for the tests; the common parameters and their values are summarized in Table 1.[6]

For all cases, we consider networks consisting of equal-sized cells, with a single base station at the center of each cell. For simplicity, the cells are square-shaped. The number of users in each cell is a random variable uniformly distributed between one and a given maximum number that varies across the cases. The locations of the users in a cell are uniformly distributed throughout the cell. For each of the cases, 50 different location scenarios are evaluated with respect to ($P3$), both distributed power allocation algorithms, and the autonomous algorithm from [8]. In all algorithms there is an update step where the policy is updated, evaluated, and compared with the previous best policy. In order to compare algorithm complexity, for each case, we will report the number of update steps performed in each algorithm averaged over the 50 location scenarios. We will also report the number of iterations needed for the distributed algorithms to converge. Here an iteration consists of a complete cycle through all cells, where during each step the appropriate base station allocates powers to its users according to *** for the first distributed algorithm, and according to (9) and (10), for the second distributed algorithm; we did not exploit any parallelism that may have been possible.

We adopt the distance loss model in [10] (Eq. (2.192) in Section 2.5 on p. 94) with log-normal fading (shadowing). Specifically, the instantaneous channel gain g between an arbitrary user and base station separated by a distance of d (in meters) is given by

$$10 \log_{10} g(d) = 10 \log_{10} \left(4 \left(\frac{\lambda_c}{4\pi d} \right)^2 \sin^2 \left(\frac{2\pi h_b h_m}{\lambda_c d} \right) \right) + \ u \text{ dB},$$

where u is a zero-mean Gaussian random variable with log standard deviation 6 dB, and where h_b and h_m are the heights (in meters) of the base station and user antennas, respectively, and $\lambda_c = 0.1579$ m ($f_c = 1.8$ GHz) is the wavelength of the transmitted signal. Note that the channel gain is not monotonically decreasing with distance.

The first two cases consist of nine cells, each with at most three users, arranged in a three-by-three array but with different cell sizes. In case 1, the base stations are 2000 m apart (typical macrocell), and in case 2, the base stations are 200 m apart (typical microcell). The maximum The average number of users in the system is the same in both cases, so the user density is greater

[6] Recall that power in dBm is given by 10 log(power in Watts/1mW).

Parameter	Value
Carrier frequency f_c	1.8 GHz
Number of cells	9
Maximum number of users per cell	3
Height of base station antenna	20 m
Height of mobile user antenna	1.5 m
Distance loss exponent	4.0
α	1.0
Log standard deviation of the log-normal random variable	6 dB
Background interference and thermal noise power I	-70 dBm
Maximum transmission power from a user	100 mW

Table 1: Parameters and their values used in the computer simulation.

in case 2 than case 1. For each scenario, we determine the number of iterations for convergence, and compare the resulting power allocation with the globally optimal power allocation obtained by solving ($P3$). As noted previously, the objective function J in ($P3$) is not convex in $\boldsymbol{p}$; the only search method that is guaranteed to find the optimal allocation is exhaustive search. Due to the "curse of dimensionality," we found that finding the optimal solution for networks of more than nine cells with three users was computationally prohibitive. Indeed the average number of comparision steps needed to find the optimal allocation was 1104282 in case 1 and 719350 in case 2. We also tested a linear array of cells (case 3) where there were six cells and at most five users per cell. In this case the average number of comparision steps needed to find the optimal allocation was 5026652.

The numerical results for the three cases are reported in Table 2. Results for both distributed algorithms plus the autonomous algorithm are reported. The first column of numbers gives the number of test scenarios (out of 50) where the optimal policy was not found by the given algorithm. The second and third columns give the average percentage difference and the maximum percentage difference from the optimal solution over all 50 scenarios, respectively. The fourth column gives the average number of iterations for the given algorithm to converge. Finally, the last column gives the average number of comparision steps executed by the algorithm.

In case 1 (2000 m between base stations), the optimal power allocations were quite diverse and ranged from all 23 users transmitting to only six users (out of a possible 17) in the whole network transmitting. For all scenarios the first heuristic distributed power allocation algorithm converged within three iterations. In particular, 38 scenarios converged in two iterations, and twelve in three. However the second distributed algorithm took an average of 19.16 iterations to converge. Similar differences can be seen in cases 2 and 3. However, this added complexity led to a small performance advantage for the second distributed algorithm. In all three cases it slightly dominated both in terms of number of cases where the optimal was found and in terms of average percentage difference from optimal.

In general, the number of users transmitting in case 2 was far fewer than in case 1, and ranged from only one out of 20 users to seven out of 18 users. This is to be expected due to the small cell size which results in higher levels of intercell interference, and thus fewer simultaneous transmissions. It

	algorithm	# not opt.	% diff.	max. diff.	avg # itn.	comp. step
case1	distrib. 1	3	0.03	0.99	2.24	41
	distrib. 2	0	0.00	0.00	19.16	95.88
	auton.	21	1.09	24.89	2.02	36.96
case2	distrib. 1	14	2.02	42.38	3.04	55.04
	distrib. 2	3	1.71	42.38	26.9	131.92
	auton.	50	62.45	98.33	2.48	45.28
case3	distrib. 1	6	0.09	2.08	2.26	41.76
	distrib. 2	0	0.00	0.00	12.8	179.08
	auton.	23	0.92	13.58	2.04	37.46

Table 2: Numerical results for computer simulation of cases 1-3.

can also be seen that the autonomous algorithm had a much higher distance from optimal than in case 1, which is again due to the higher levels of intercell interference, making its greedy assignments much less efficient.

It appears that when it is optimal for only a few users transmit, the chance that either distributed algorithm converges to the optimal allocation is reduced because the objective function is highly sensitive to each user. For example, in one of the two scenarios of case 2 where the percentage difference for the first distributed algorithms from optimal was greater than 50% only two users transmitted under both the optimal allocation and the distributed. However, had the cells in which the transmitting users were located gone first in each step of each iteration, then the distributed algorithm would have converged to the optimal allocation. Therefore, when it is optimal for only a small number of users to transmit, then the manner in which the base stations are ordered within each iteration becomes critical. Indeed, when the the maximum transmission power is decreased by a factor of ten in case 2, or the background interference and thermal noise power is increased by a factor of ten, it becomes optimal for several users to transmit simultaneously in each cell, and the percentage of scenarios for which the two distributed algorithms converged to a power allocation that was the solution of $(P3)$ was 86% and 98% for each distributed algorithm, respectively.

In case 4 we investigated a bigger example, still with 50 different scenarios. It was not possible to find the optimal solution for this case but we did compare the two distributed algorithms and the autonomous algorithm. The example consists of twelve cells, each with at most ten users, arranged in a three-by-four array but with base stations that are 200 m apart. Here we use the second distributed algorithm as the basis for comparison. It took on average 8137 comparison steps to converge whereas the first distributed algorithm took 210 steps and the autonomous algorithm took 163 comparison steps. There were 17 cases where the two distributed algorithms differed with the average percentage difference over all cases being 0.39% and the maximum difference being 7.17%. The autonomous algorithm found a different solution in all 50 cases. The average percentage difference was 54.93% and the maximum percentage difference was 97.17%.

6 Final remarks

In this paper, we provided a distributed version of the algorithm, where base stations make use of the available knowledge of the states of the users in neighboring cells to control the transmission power and spreading gains of the users. Since the optimal spreading gain control (Proposition 1) is effectively distributed, the main effort centered on designing a distributed power control algorithm that converges to the globally optimal (or near optimal) solution of $(P2)$. The power control algorithm given in Section 4 is simple and converges to a Nash equilibrium; simulations show that it frequently converges rapidly to the globally optimal power allocation.

Since wireless communication links (radio channels) typically have low bandwidths (compared to their wired counterparts) and are of time-varying quality, an important design goal is to develop resource allocation algorithms that squeeze the most out of the scarce radio communication resources. One way to do this is to take advantage of the times when channel qualities are good, varying the transmission powers and spreading gains of the users accordingly. Here the transmission powers and spreading gains were controlled jointly to maximize aggregate throughput, subject to peak power constraints. It would be of interest in future research to suitably exploit the results of this paper to develop distributed resource allocation algorithms providing short-term performance guarantees to delay tolerant users, without undue degradation to aggregate throughput.

References

[1] S. V. Hanly, "An algorithm for combined cell-site selection and power control to maximize cellular spread spectrum capacity," IEEE J. Select. Areas Commun., vol. 13, no. 7, pp. 1332–1340, 1995.

[2] R. Yates and C.-Y. Huang, "Integrated power control and base station assignment," IEEE Trans. Veh. Tech., vol. 44, no. 3, pp. 638-644, 1995.

[3] S. Grandhi, J. Zander, and R. Yates, "Constrained power control," Internat. J. Wireless Person. Commun., vol. 1, no. 4, 1995.

[4] D. Mitra and J. A. Morrison, "A distributed power control algorithm for bursty transmissions on cellular spread spectrum wireless networks," in Wireless Information Networks, J. M. Holtzman, Ed., Kluwer, Boston, MA, 1996.

[5] D. Kim, "Rate-regulated power control for supporting flexible transmission in future CDMA mobile networks," IEEE J. Select. Areas Commun., vol. 17, no. 5, pp. 968-977, 1999.

[6] H. Ji and C.-Y. Huang, "Non-cooperative uplink power control in cellular radio systems," Wireless Networks, vol. 4, no. 3, pp. 233–240, 1998.

[7] N. Feng, N. B. Mandayam, and D. J. Goodman, "Joint power and rate optimization for wireless data services based on utility functions," in Proc. Conf. Inform. Sci. and Sys. (CISS'99), Johns Hopkins University, Baltimore, MD, 1999.

[8] S.-J. Oh and K. M. Wasserman, "Optimal resource allocation in wideband CDMA networks," preprint, 1999. (http://www.eecs.umich.edu/~wass/wcdma.ps)

[9] S. Kumar and S. Nanda, "High data rate packet communications for cellular networks using CDMA: Algorithms and performance," IEEE J. Select. Areas Commun., vol. 17, no. 4, pp. 472-492, 1999.

[10] G. L. Stüber, Principles of Mobile Communication. Kluwer Academic Publishers, Boston, MA, 1997.