Distributed Optimal Rate Assignments for CDMA-Based Wireless Networks

Jennifer Price, Tara Javidi
{pricej,tara}@ee.washington.edu
Dept of EE, University of Washington
Seattle WA, 98195-2500

UWEE Technical Report
Number UWEETR-2004-0022
August 2004

Department of Electrical Engineering
University of Washington
Box 352500
Seattle, Washington 98195-2500
PHN: (206) 543-2150
FAX: (206) 543-3842
URL: http://www.ee.washington.edu

Distributed Optimal Rate Assignments for CDMA-Based Wireless Networks

Jennifer Price, Tara Javidi
{pricej,tara}@ee.washington.edu
Dept of EE, University of Washington
Seattle WA, 98195-2500

University of Washington, Dept. of EE, UWEETR-2004-0022
August 2004

Abstract

Increased demand for wireless broadband data services has led to a reexamination of system design, particularly with respect to optimal resource allocation. In this paper, we examine this optimal resource allocation as a rate assignment problem in terms of an optimization framework. We then develop a distributed algorithm for solving this problem, and examine its behavior under time-varying conditions. Finally, we propose modifications to this algorithm to mitigate its undesirable dynamic behaviors.

1 Introduction

As the demand for broadband wireless data services grows, it becomes necessary to re-examine system design with respect to resource allocation. Due to the inherent differences between voice and data traffic (in both the QOS requirements and statistical behavior), the design of future broadband networks must be examined in a new context, particularly with respect to optimal resource allocation.

We choose to examine this as a rate assignment problem using an optimization framework. Conceptually, this can be viewed as a natural extension of existing work on congestion control in wireless TCP (see [7], [8], [10], [11], [13]). Specifically, we focus on a wideband CDMA network with arbitrary but known layout and variable transmission rates, similar to the CDMA2000 and 1xEVDO systems. Although work has been done on rate-assignment in WCDMA and CDMA2000 systems ([3], [4]), until recently no work proposed a decentralized, non-uniform rate assignment.

The rest of this paper is organized as follows. Section 2 develops a common interference model for CDMA, a definition of feasible rate assignments, and the formal optimization problem. Section 3 describes the formulation of a distributed algorithm for rate assignment. Section 4 introduces several practical considerations including non-homogenous demand and dynamic behavior, and proposes modifications to improve algorithm performance. Finally, Section 5 presents our conclusions and areas of future work.

2 System Model

2.1 CDMA Interference Model

We use the following notation. There are a total of N mobiles and L sectors. The *tracking base* for mobile i (denoted by $b(i)$) is the base to which the mobile is connected and which transmits power control signals to the mobile. For simplicity, we assume that each mobile is tracked by exactly one base. We further assume that if $l = b(i)$, then $g_{il} > \epsilon$, where g_{il} is the channel gain (this includes both path and antenna gain).

P_i is the transmitted power for user i, and x_i is the transmitted rate for user i. W is the chip bandwidth, N_0 is the thermal noise density. We express the transmission rate of each mobile as a multiple of the wireless pilot rate, R_b (i.e. $x_i = R_b\alpha_i$). The *spreading gain* for mobile i is defined as $s_i = \frac{W}{R_b\alpha_i}$. Note that $0 \leq R_b\alpha_i \leq W$. We introduce the following change of variable for simplification:

Definition 1. *We define the effective rate of mobile i to be $r_i = \frac{x_i}{W+\gamma x_i} = \frac{R_b \alpha_i}{W+\gamma R_b \alpha_i}$.*

Now, consider mobile i which is tracked by sector $l = b(i)$. The signal to noise ratio of mobile i at the base station l can be written as

$$SINR^l(i) = \frac{s_i P_i g_{il}}{N_0 W + \sum_{j \neq i} P_j g_{jl}} \tag{1}$$

Definition 2. *The rise over thermal (ROT) is defined as*

$$Z_l = \frac{\sum_{i=1}^{N} P_i g_{il}}{N_0 W} \tag{2}$$

and indicates the ratio between the total power received from all the mobiles at the base station l and the thermal noise [1].

2.2 Rate Feasibility Region

We choose to examine the rate assignment problem in terms of an optimization framework. In other words, we seek to optimize some function of the rate assignment, subject to the interference constraints of the wireless medium. In order to this, we must establish a feasible region of rate vectors. It has been shown in [5], [6], [14] that such a feasibility region can be defined as follows.

Definition 3. *A vector of rates $(x_1 \ldots, x_N)$ belongs to the feasible region Δ if and only if the corresponding vector of effective rates $(r_1, \ldots, r_N)$ satisfies the following condition:*

LC1. $\sum_{i=1}^{N} r_i \frac{g_{il}}{g_{ib(i)}} \leq \frac{K}{\gamma(1+K)}$

Here, γ is a pre-specified value such that satisfying $SINR \geq \gamma$ guarantees an acceptable BER [15]. Note that Condition LC1 is sufficient (see [5]) to ensure that $Z_l \leq K$ [1], [14] which, in turn, is sufficient (see [5], [14]) to ensure that the instantaneous transmit power for each user is small [1], [5].

2.3 Rate Assignment: Throughput versus Fairness

The feasibility region described in Section 2.2 contains a possibly infinite number of rate vectors. Of these vectors, we wish to choose the vector that is optimal with respect to some performance criteria, specified by an objective function. While the authors have considered various criteria such as maximum total throughput, max-min fairness, and proportional fairness [5], [6], here we focus on a weighted proportional fair solution. It has been shown in [8] that the function $U(\cdot) = log(\cdot)$ results in a proportional fair solution. Multiplying this utility function by a weight w_i for each user results in a weighted proportional fair solution. Combining this with our feasibility region, we define the following problem:

P. Find the rate vector $(x_1, \ldots, x_N)$ such that the corresponding vector of effective rates $(r_1, \ldots r_N)$ solves the following:

$$\max_{r \in \Delta} \sum_{i=1}^{N} w_i U(r_i)$$

where w_i is the weight associated with user i.

In general, we would like to have a strictly increasing and concave utility function, which can be achieved by restricting our attention to rates $0 \leq r_i < \frac{1}{2\gamma}$.

The use of weights allows us to distinguish several classes of users by weighting their utility functions. In other words, we may wish to declare subsets of users as more or less "important," and allocate the bandwidth appropriately. This classification could be based on any number of factors, including power limitations, type of data, or packet backlog. This essentially allows the incorporation of a priority for each user. Additionally, the author in [9] notes that in a weighted optimization problem, choosing integral weights is equivalent to user i representing itself as w_i identical sub-users. If we interpret these weights as a charge which is proportional to the bandwidth received, we can say that the solution to Problem P results in the rates per unit charge being proportionally fair [7].

3 Distributed Algorithm Design

The authors in [5] and [14] have developed a distributed rate assignment algorithm that *approximates* the solution to problem P. Conceptually, this algorithm is based on a simple auctioning mechanism [2], and consists of two components:

1. Each base station produces a regulating/coordinating signal that indicates the level of interference at each sector (price announcement).

2. Each mobile reacts to the levels of interference (indicated by the base station coordination signals) by adjusting its rate to maximize the profit (utility minus cost).

It can be shown (see [5], [14]) that this results in a regulation regime which aligns the selfish behavior at the mobiles with the social welfare. Here we describe a practically implementable modification of this algorithm in detail.

Base Algorithm

Each base, l, computes its price, μ_l, to align the "selfish" mobile strategies with social welfare [12]. This is accomplished when μ_l is aligned with the Lagrangian multipliers from Problem P. Hence, the base algorithm uses a gradient projection method to compute μ_l [11]. We approximate this gradient method by the following equation:

$$\Delta\mu_l = \begin{cases} \beta(Z_l - K) & \text{if } \mu_l > 0 \\ \beta[Z_l - K]^+ & \text{if } \mu_l = 0 \end{cases} \tag{3}$$

This allows the bases to announce their prices according to $\mu_l^t = \Delta\mu_l + \mu_l^{t-1}$.

Mobile Algorithm

Each mobile computes its (selfishly) optimal rate at any computation epoch. In order to do so, each mobile needs to compute its weighted price (proportional to the sum of its contribution to the ROT at each sector). Thus, for a given vector $\underline{\mu} = (\mu_1, \ldots, \mu_L)$ of prices declared by the base stations, each mobile calculates its own effective rate at time t, r_i, such that:

$$r_i^t = \arg\max_{r_i}(U(r_i) - r_i p_i^t) \tag{4}$$

where $p_i^t = \sum_{l=1}^{L} \frac{g_{il}}{g_{ib(i)}}\mu_l$ is the total price for mobile i at time t. Note that the economic interpretation of this algorithm is that for any price p_i, each mobile maximizes its profit (utility minus cost).

In calculating their own selfish optimal rate, it might seem that each mobile requires full knowledge of the channel, including its gain to all bases, in order to use Eqn (2). In [5] and [14], the authors show that there exists a practical solution to this problem using the CDMA pilot signal, PS, and a pricing pilot signal, PPS. The pricing pilot signal is identical to the pilot signal with the only difference that it is transmitted with a power level proportional to the base price, μ_l. The result is that p_i can be calculated as follows:

$$p_i = \sum_{l=1}^{L} \frac{g_{il}}{g_{ib(i)}}\mu_l = \frac{E_{TR}^{PPS}}{E_T^P(b(i))} \tag{5}$$

where $E_{TR}^{PPS} \propto P_t^P \sum_{l=1}^{L} g_{il}\mu_l$ denotes the total PPS energy received by mobile i, and $E_T^P(b(i)) \propto P_t^P g_{ib(i)}$ is the PS energy received by mobile i from its tracking sector $b(i)$.

3.1 Implementation and Power Control

One of the main difficulties in implementing these algorithms is the issue of power control for the mobiles. Since the prices announced by the bases rely on the ROT (which, in turn, relies on the transmit powers), the algorithm requires "perfect" power control. In real scenarios, this is accomplished by running a power control loop at a time scale much faster than that of the algorithm (similar to that of a voice system [15]). The mobiles will simply adjust their transmit power up or down by a value ΔP in response to a signal from their tracking sector. When varying their rates, the mobiles also attempt to control their own transmit power so as to account for discontinuities and satisfy Condition C3 from Section 2.2. This is done by simply matching any proportional increase or decrease in rate with a proportional increase or decrease in transmit power.

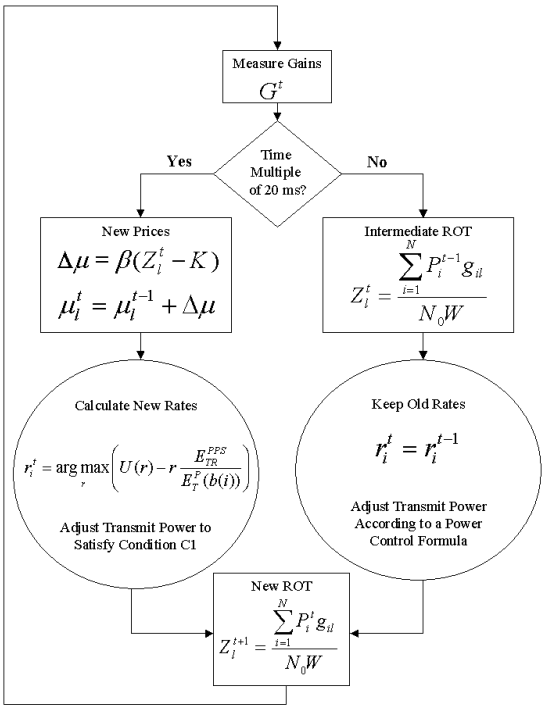

Figure 1: Flowchart of Algorithm Operation

Figure 1 shows the operation and interaction of the mobile and base algorithms. The base stations measure updated channel information every 5 ms - i.e., the loop shown in the flowchart is run every 5 ms. The bases calculate new prices every 20 ms, and immediately broadcast them through the use of the Pricing Pilot Signal. The mobiles respond to the receipt of new prices by calculating new rates. The mobiles continually adjust their powers using the power control loop discussed earlier. The squares represent components of the base algorithm, while the circles represent components of the mobile algorithm.

4 Practical Issues and Heuristic Modifications

In order to examine the behavior of this algorithm, simulations were run using a symmetric layout of four base stations (each 2500m apart) and 20 mobiles. The simulations use a cost-231 propagation model at 1.9 GHz between the mobiles and bases. The values for γ and K are 4 dB and 6 dB. The chip bandwidth W is 1.2 MHz, and the pilot rate R_b is 4.8 Kbps. For simulations involving mobility, a subset of five mobiles were allowed to move using a model of constant speed (120 km/hr) and random direction. As discussed in the previous section, bases obtain updated channel information and run the power control loop every 5 ms using a ΔP of .25 dBm. The *Base Algorithm* and is run every 20 ms, and the mobiles respond to new prices by running the *Mobile Algorithm* to generate new rates. Figure 2 shows a contour plot illustrating the equilibrium point of the distributed rate assignment algorithm using a proportional fair scheme.

4.1 Non-Homogenous Demand

Although we focus on a rate-based system, actual implementations are packet-based. We can reconcile these two ideas by thinking of rate as the number of packets transmitted during a unit of time. Obviously, if a mobile has a limited number of packets to send, it does not benefit from the allocation of additional bandwidth. In the case where a mobile has no packets to send, we can think of this as the removal of a user from the system. In the problem formulation presented in Section 2.3, each mobile has an identical, strictly increasing utility function. Thus, the algorithm does no distinguish between users that have traffic to send and those that do not. The result is that the “leftover” bandwidth from users without packets to send is wasted, and then network becomes underutilized.

4.1.1 Proposal: User-Specific Utility Functions

One way to model this phenomenon is to introduce the concept of non-homogenous demand; in other words, we incorporate how much bandwidth each user needs into the utility functions. This results in the following utility functions:

$$U_i^1(r_i) = \begin{cases} log(\frac{Wr_i}{1-\gamma r_i}) & \text{if} \quad r_i < r_i^{\max} \\ log(\frac{Wr_i^{\max}}{1-\gamma r_i^{max}}) & \text{else} \end{cases} \tag{6}$$

The result is that each user has a unique utility function which is no longer strictly increasing over the allowable interval and saturates at a maximum transmission rate, r_i^{max}. We can simply substitute the new utility functions using r_i^{max} to indicate each users' demand for bandwidth, and pose the following modified problem:

P2. Find the rate vector $(x_1, \ldots, x_N)$ such that the corresponding vector of effective rates $(r_1, \ldots r_N)$ solves the following:

$$\max_{r \in \Delta} \sum_{i=1}^{N} w_i U_i^1(r_i)$$

For example, r_i^{max} can be chosen as the total number of remaining packets divided by units of time. While this does not exactly capture the dynamics of time-varying non-homogenous demand, it does allow us to use a natural extension of the original distributed algorithm, as follows:

Modified Base Algorithm

Identical to the original Base Algorithm described in Section 3.

Modified Mobile Algorithm

Same as the original Mobile Algorithm described in Section 3, but use $U_i^1(r_i)$ instead of $U(r_i)$.

Although this distributed algorithm is similar to the one described in Section 3, the equilibrium point of this algorithm changes. In order to more carefully examine what happens when a user is maximum-rate limited in the multi-sector case, we assume user N is under-utilizing its bandwidth and that the values r_i^{max} are independent of time. The remaining wireless users will change their rates from α_i to $\alpha_i + \Delta\alpha_i$. The condition for optimality of the $\Delta\alpha_i$'s can be written as:

$$\sum_{i=1}^{N} U_i' \left(R_b(\alpha_i + \Delta\alpha_i)\right) \left(\Delta\alpha_i' - \Delta\alpha_i\right) \leq 0 \ \ [7]$$

where $\Delta\alpha_i'$ is any other value. This reduces to

$$\sum_{i=1}^{N} \frac{\Delta\alpha_i' - \Delta\alpha_i}{\alpha_i + \Delta\alpha_i} \leq 0$$

In [8], a vector $\underline{x}$ is defined to be proportional fair if it satisfies the relation $\sum_{i=1}^{M} \frac{x_i' - x_i}{x_i} \leq 0$. From this definition, we can see that the excess bandwidth is not split in a strictly proportional fair manner. The change in rates are, however, proportional fair with respect to the total rates. It is also clear that the total rate vectors for the remaining $N-1$ users are proportional fair, since $U_N'(\alpha_N) = 0$.

Figure 3 shows the contour plot illustrating the equilibrium point of the rate assignment algorithm when users have non-homogenous demand, reflected in their utility functions. As expected, we see that users who cannot fully utilize their allocated bandwidth are assigned lower rates. This excess bandwidth is then distributed among the remaining users, still resulting in a proportional fair rate allocation among users who do no have maximum rate limitations.

4.2 Dynamic Behavior

While it is important to note that the above algorithm exhibits "good" static behavior under time-invariant conditions, the more interesting issue is the behavior of the algorithm when the system parameters, e.g. channel gains, sector assignment, etc. vary with time. It is shown in [5] and [14] that the proposed distributed algorithm converges to an equilibrium when the network topology and gains are static. However, dynamic variations can cause significant performance degradation and may need to be addressed. Through extensive simulations and observation (see [5], [6],

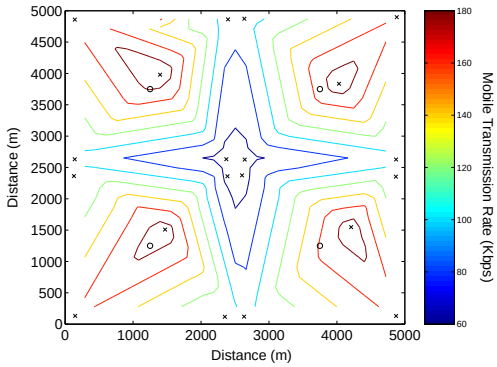

Figure 2: Contour Plot of Equilibrium Rate Assignments for a Symmetric Layout

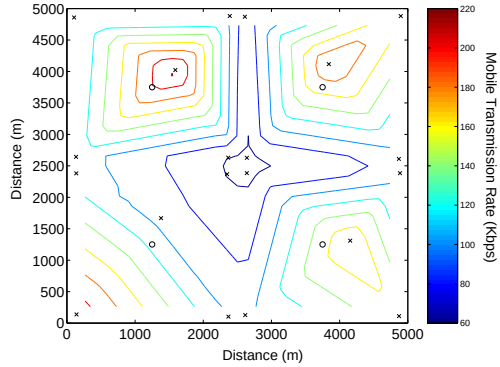

Figure 3: Contour Plot of Equilibrium Rate Assignments for Non-Homogenous Demand

[14]), we have classified these dynamic variations into three types: fast and random variations, sudden but quasi-static topological changes, and persistent dynamic variations.

Fast and Random Variations

Fast and random variations include the use of imperfect power control and mobility. These dynamics tend to introduce a level of uncertainty around the equilibrium point, but only cause minor performance degradation. As seen in Figures 4 and 5, the system performance remains close to the predicted equilibrium point even when considering mobile speeds on the order of 120 km/hr. As such, we believe that mitigating these effects would provide very little benefit, and would not justify the trade-off in computational complexity.

Persistent Dynamic Variations

Persistent dynamic variations in network setting occur continually over time. Unlike random environment changes, however, persistent dynamic variations do cause significant performance degradation and require some level of compensation. When considering these types of dynamics, the transient mode is the dominant mode. In other words, the time constants of significant change in network setting and the distributed algorithm are comparable. In the case of shadowing, the times scales of change for the gains $g_{il}(t)$ are at the same order as the time constants of the distributed algorithm. In Figures 4 and 5, we see the effect of shadowing on the algorithm performance. Although shadowing does cause fairly large variations in both rate and ROT levels, these variations do tend to occur around the desired equilibrium point when K is chosen conservatively (low level of interference). The impact of the choice of K on stability remains an interesting topic of further research.

As previously mentioned, changes in user demand may also fall under the category of persistent dynamical variations. This is the case when the time scales of change for the variables $r_i^{max}(t)$ are at the same order as the time constants of the distributed algorithm. In fact, given the behavior of current TCP protocols, it is much more likely that users' demands will continually change over time. It remains to be seen how to best integrate such changes over time into the algorithm design.

Sudden but Quasi-Static Topological Changes

Sudden topological changes can be classified as quasi-static: that is, they cause sudden and significant changes but tend to occur sparsely over time. Examples of such changes include the addition and removal of users from the system. Depending on the behavior of upper layer protocols, specifically the transport layer, changes in user demand may (or may not) be classified as quasi-static. This is the case when the variables r_i^{max} change in a quasi-static manner, as described in Section 4.1. These types of changes, while occurring infrequently, tend to cause a significant spike in ROT at the affected base, violating the constraints C1 and C2. In other words, the dynamic variation in the setting negatively impacts the performance by:

1. a sharp increase in ROT above the threshold, K, which results in a period of outage, and

2. a "low performance" transient period with increased bit error rate (BER).

In Figures 4 and 5, we see the effect of adding a new user at time $t = 100ms$. The spike in ROT that accompanies the addition of a new user violates system constraints, causing unacceptably high BER. In addition, we see that the new user sees not only a high initial transmit rate, but a large back-off in rate and a fairly slow convergence time. In order to mitigate these effects, we propose modifications to the mobile algorithm that will:

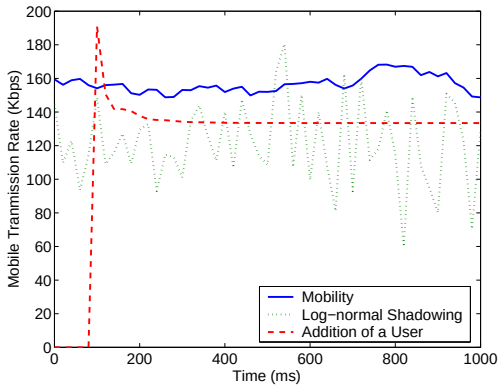

Figure 4: Rate Assignment for an Individual User in the Presence of Dynamic Variations

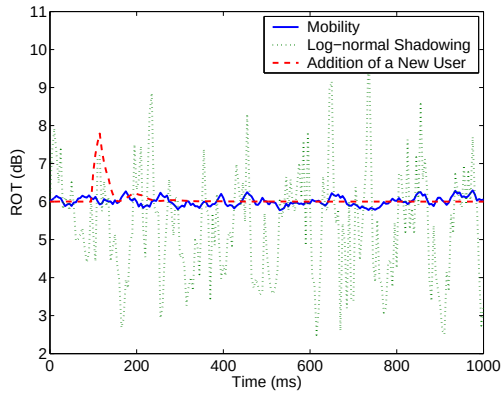

Figure 5: ROT Levels in the Presence of Dynamic Variations ($ROT > 6dB$ yields unacceptable BER)

1. mitigate the overshoot of ROT during the transient period, and

2. increase the rate of convergence, thereby shortening the transient period.

These modifications and their performance are described below.

4.2.1 Proposal: Compensation for Quasi-Static Changes

Based on the previous section's examples, it is clear that the existing basic algorithm can exhibit undesirable behavior under various conditions. In order to improve this algorithm, we propose the following heuristic mobile algorithms. These algorithms are designed specifically with sudden quasi-static topological changes in mind. Specifically, we seek to reduce the overshoot in ROT, and increase the rate of convergence after these changes occur.

Slow-Start

One way to reduce the spike in ROT seen in the previous section is to use a slow-start mechanism for new users, accomplished by imposing a maximum increase in rate. The mobile algorithm becomes:

$$r_i = \min[\arg\max_r (U(r) - rp_i),\ r_i^{t-1} + \Delta r] \tag{7}$$

where Δr is the maximum allowable rate increase.

This algorithm forces new users to slowly increase their rates, allowing the prices to catch up with the added traffic rate.

Look-Ahead

The undesirable transient behavior of the basic algorithm has an important economic interpretation: in a CDMA system with a small number of users, the user may not be a price-taker agent (i.e. the mobile can affect the price significantly by varying its rate). At each computation epoch, the original algorithm chooses its myopic optimal rate, ignoring its own effect on the future price. The idea behind a look-ahead algorithm is to have the mobile anticipate the change in price based on its new rate assignment.

The initial rate, r_i', is calculated exactly as in the basic algorithm. Using this change in rate and the gain information included in the CDMA pricing signal, the mobile can estimate the resulting ROT at its tracking sector ($Z'_{b(i)}$). The mobile then calculates its actual rate, r_i^t, using the formula:

$$r_i^t = \begin{cases} \arg\max_{r_i}(U(r_i) - r_i p_i') & \text{if } |Z'_{b(i)} - K| > 0\ dB \\ r_i' & \text{else} \end{cases} \tag{8}$$

where $p_i' = p_i^t + \beta(Z'_{b(i)} - K)$.

Figures 6 and 7 compare the performance of the basic, slow start, and look-ahead algorithms when a new user is added at time $t = 100ms$. Figure 6 shows the rate assignments for the new user, while Figure 7 shows the ROT at the base to which the new user is added. As expected, the slow start algorithm is the most conservative and takes the longest to converge. The look-ahead algorithm performs the best, as it eliminates the spikes in rate and ROT while still converging faster than the original algorithm. Understanding the exact nature of these algorithms requires further investigation.

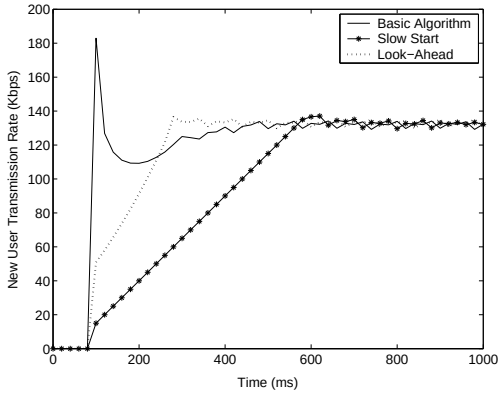

Figure 6: Rate Assignments for an Added User Under Various Algorithms

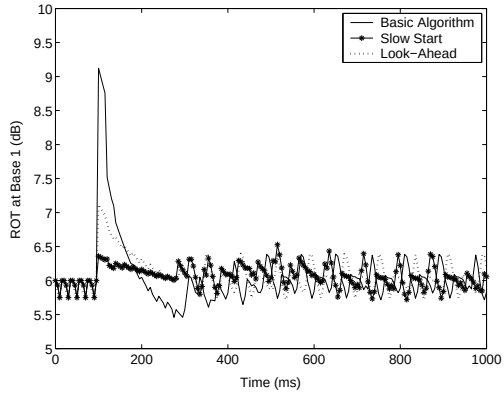

Figure 7: ROT Levels for Addition of a New User Under Various Algorithms ($ROT > 6dB$ yields unacceptable BER)

5 Conclusions

In this paper, we address the issue of decentralized rate allocation in a wideband CDMA system serving users with variable rates. We formulate the rate assignment problem as a global optimization problem, which accommodates several classes of users. We then describe an existing distributed algorithm for rate assignment, and introduce utility functions that allow for non-homogenous demand without modifying the existing algorithm. Finally, we study the performance of this algorithm in dynamic settings, and propose modifications to mitigate undesirable behaviors.

While the modifications presented here do improve the algorithm performance, there is room for further investigation. Obviously, we wish to further address the algorithm performance in the presence of significant dynamic variations, particulary changing user demand. In order to compensate for these variations, an accurate model of the algorithm's dynamic behavior is necessary. This will almost certainly allow for a better understanding of the algorithms behavior, and the best way to compensate for changes. In addition, the question of how to integrate changing user demand must be addressed. A cross-layer (MAC and transport) approach to joint rate and congestion control may be an appropriate way to handle changing user demand. The key will be to develop an analytic model that accurately describes the system under realistic conditions.

Acknowledgments

This work was supported in part by the National Science Foundation ADVANCE Cooperative Agreement No. SBE-0123552. The authors would also like to thank Dr. C. Lott and Dr. D. Zhang at Qualcomm Corporate Research and Development for their thoughtful comments and suggestions.

References

[1] 3rd Generation Partnership Project 2 (3GPP2). CDMA2000 high rate packet data air interface specification. Technical Report CS20024, www.3gpp2.org, 2001.

[2] K. J. Arrow and F. H. Hahn. *General Competitive Analysis*. Holden-Day Inc., 1971.

[3] S. Chakravarty, R. Pankaj, and E. Esteves. An algorithm for reverse traffic channel rate contol for CDMA2000 high rate packet data systems. In *Proceedings of IEEE Globecom*, San Antonio, TX, 2001.

[4] E. Esteves. On the reverse link performance of CDMA2000 high rate packet data systems. In *Proceedings of IEEE International Communication Conference*, 2002.

[5] T. Javidi. Decentralized rate assignments in a multi-sector cdma network. In *Global Telecommunications Conference*, pages 65–69, 2003.

[6] T. Javidi and J. Price. Decentralized rate assignments in a multi-sector cdma network. In *In Preparation for Publication*.

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

[8] F. Kelly. Mathematical modelling of the internet. In *Proceedings of the Fourth International Congress on Industrial and Applied Mathematics*, pages 105–116, 1999.

[9] F.P. Kelly, A.K. Maulloo, and D.K.H. Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.

[10] R.J. La and V. Anantharam. Charge-sensitive TCP and rate control in the Internet. *IEEE Communications Magazine*, pages 74–80, Oct. 2003.

[11] S. H. Low and D. E. Lapsley. Optimization flow control, i: Basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, 7(6):861–874, Dec 1999.

[12] S. H. Low, F. Paganini, and J.C. Doyle. Internet congestion control. *IEEE Control Systems Magazine*, February 2002.

[13] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):555–567, Oct 2000.

[14] J. Price and T. Javidi. Decentralized and fair rate control in a multi-sector CDMA system. In *Proceedings of the Wireless Communications and Networking Conference*, 2004.

[15] R. D. Yates. A framework for uplink power control in cellular radio systems. *IEEE Journal on Selected Areas in Communications*, 13(7):1341–1347, Sept 1995.