Decentralized Rate Assignments in a Multi-Sector CDMA Network

Tara Javidi
Electrical Engineering, University of Washington, Seattle
tara@ee.washington.edu

***Abstract*— We consider a wideband CDMA network with arbitrary but known layout of sectors and base-stations, and with variable mobile rates. In this paper, we investigate the issue of rate assignment in a CDMA network. We show that in a broadband wireless data system based on a wideband CDMA technology, there exists an optimal decentralized rate assignment. We show how this method is related to the traditional window or rate based flow control mechanisms widely used in TCP/IP networks.**

I. Introduction

As the applications of wireless networks (such as cellular systems) spreads to various fields, the demand for wireless broadband data (such as video over wireless) services increases. Due to inherent difference between statistical behavior and Quality of Service (QoS) requirements of a data connection with those of a voice user, new challenges and problems in the design of future generations wireless systems arise. More specifically, in such systems, where throughput and fairness become essential elements of QoS and service provisioning the need to revisit the question of optimal resource allocation is crucial. We believe that it is of extreme importance for the wireless broadband data networks to have similar characteristics as those of the modern IP networks. In particular, it is important for such networks to respond to randomly fluctuating demands and failures by adapting rates in a scalable manner, essentially similar to that of the flow and congestion control mechanisms in TCP.

On the other hand, the wireless medium adds further challenges which need a special care. In particular, the issue of medium access control (MAC layer) becomes a crucial issue in designing the airlink in any wireless system with large coverage area and seamless service. Unlike local area networks, where simple random access schemes provide sufficient utility, in systems with large range and coverage (wide area broadband systems) more sophisticated MAC layers are deployed to address the challenges of bandwidth sharing by multiple users. Such technologies include various spectrum spreading techniques such as wideband CDMA and OFDMA. We do not intend, in this paper, to compare these systems, but rather provide a methodology which allow us to address issues of fairness, optimal channel allocation, dynamic rate control, stability, scheduling complexity, and decentralization of information and control in such systems. We believe such a methodology is essential to comparison studies of the MAC layer technologies.

In this paper, we focus on a wideband CDMA network with an arbitrary but known layout of sectors and base-stations, serving variable mobile rate, a system similar to the recently deployed commercial CDMA2000 1xEV-DO systems in South Korea and parts of US. We believe that this work sheds light on similar issues arising in an OFDM-based systems in a different context. We believe a similar analysis (from a network perspective) of OFDM-based systems is required in order to understand the pros and cons of each of these spread-spectrum techniques. Here it suffices to emphasize that the main difference between OFDM-based systems and wideband CDMA is in how these technologies tackle the fundamental trade-off between instantaneous (myopic) optimal allocation of bandwidth (static utility) on the one hand and the complexity, stability, and sensitivity of such allocation mechanism under a fast changing and uncertain dynamics on the other.

In this paper we address the issue of decentralized rate assignment in a multi-sector wideband CDMA networks. The rate assignment problem in WCDMA and CDMA2000 systems has recently received some attention (see [3], [4]), but none of these studies propose a decentralized non-uniform rate assignment. In Section III, using a common interference model for CDMA systems, we investigate the existence of a feasible rate allocation in such a power controlled cellular systems. Understanding the role power control and its effects are crucial in determining a constraint describing the feasibility of a rate allocation vector, i.e. a feasibility region. In Section IV, we formulate an optimal rate allocation problem. We discuss how such a problem can be formulated to achieve high overall throughput or fairness. We propose a linearized constraint to reduce the complexity of the optimization problem. In Section V, we show that under the linear feasibility region finding an optimal rate allocation reduces to a general utility optimization framework. In other words, in Section V, we propose two sets of algorithms: 1) the flow control algorithm implemented at each mobile to react to increased interference, and 2) the base-stations' regulation algorithms which use the information available at the station to manage interference. The flow control algorithm at mobiles uses a combination of its desirable QoS (rate utility) and interference indication signals from the base to self-tune its data flow. This is similar to rate based flow control implemented in TCP. The base-station's regulation algorithm, on the other hand, acts similarly to an active queue management algorithm in that it feeds back to the mobile terminals a signal indicating the level of interference and noise. In Section VI we provide our conclusion and future work.

II. Interference Model for CDMA Systems

We use the following notation

There are a total of N mobiles and L sectors.

The *tracking sector* for mobile i is the sector to which the mobile is connected which also transmits power control signals to the mobile (tracking the mobile) and is denoted by $b(i)$.

$M_l, l = 1, \ldots, L$ is the set of mobiles which are being tracked by sector l, i.e. $i \in M_l$ iff $l = b(i)$. For simplicity, we assume that $M_l, l = 1, \ldots, L$ are mutually disjoint.

The channel power gain from mobile i to sector l is denoted by $g_{i,l}$. $g_{i,l}$ incorporates both path gain and antenna gain. We, further, assume that if $i \in M_l$ then $g_{i,l} > \epsilon$

W is the chip bandwidth (in Hz).

P_i is the transmitted power for user i, and α_i is the transmission rate for mobile i. Note that $0 \leq \alpha_i \leq W$, where $\alpha_i = W$ corresponds to no spreading. Let's define spreading gain $s_i = \frac{W}{\alpha_i}$

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

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

where N_0 is the thermal noise density.

III. Feasible Rate-Power pairs

We say a vector of rates $(\alpha_1, \ldots, \alpha_N)$ is a feasible solution if there exist a vector $(P_1, \ldots, P_N)$ such that the following conditions are satisfied

C1. $\frac{s_i P_i g_{il}}{N_0 W + \sum_{j \neq i} P_j g_{jl}} = \gamma$, for every $i \in M_l$, where γ is a pre-specified value

C2. $0 \leq P_i \leq P_i^{\max}$

In other words, in order to establish feasibility of a vector of rates $(\alpha_1, \ldots, \alpha_N)$ we need to first solve equation (1), to calculate the appropriate power vector, then establish the validity of condition C2. It is generally more desirable to have a feasibility region which can be constructed independent of power vectors. In other words, we seek to reduce the dependency of conditions C1-C2 on the power vectors. We need the following definition to establish a set of conditions which are independent of power vector $\underline{P}$.

Definition 1: For each user i, we define the quantity *effective rate* to be $r_i = \frac{1}{\gamma + s_i} = \frac{\alpha_i}{\gamma \alpha_i + W}$.

Definition 2: We define the *normalized rate* of user i tracked by sector k to be given as $R_{ik} = \frac{r_i \Psi_{ik}}{g_{ik}}$. And the matrix of normalized rates to be defined as matrix $\mathbf{R} = [R_{ik}]$ of dimension $N \times L$.

Define the gain matrix $\mathbf{G}$ of dimension $L \times N$ such that $G_{li} = g_{il}$. Let $\underline{\mathbf{1}}_M$ be a vector of dimension M whose elements are all 1. Now using the above definition, we provide the following theorem.

Theorem 1: A vector of rates $\underline{\alpha} = (\alpha_1, \ldots, \alpha_N)$ is *a feasible rate assignment* iff the vector of effective rates $\underline{\mathbf{r}} = (r_1, \ldots, r_N)$, where $r_i = \frac{\alpha_i}{\gamma \alpha_i + W}$, satisfies conditions C1':

C1'. $\mathbf{R}(I - \gamma \mathbf{G}\mathbf{R})^{-1} \underline{\mathbf{1}}_L \leq \frac{1}{\gamma N_0 W} \underline{\mathbf{P}}^{\max}$,

where $\underline{\mathbf{P}}^{\max}$ is a vector of size N, with elements $P_i^{\max}$. We denote by Δ the *feasibility region*, i.e. the set of all rate vectors $\underline{\alpha}$ which are feasible rate assignment.

Confirming the validity of Condition C1' requires computationally complex operations. In the remainder of this paper, we introduce an alternative notion which gives rise to conditions with lower complexity.

Definition 3: The ratio between the total power received from all mobiles at the base station l and the thermal noise is called *rise over thermal* (ROT) and is denoted by Z_l.

Using the above notion, the transmitted power of mobile i is such that by

$$P_i = \frac{\gamma N_0 W (1 + Z_l)}{g_{ib(i)}(\gamma + s_i)} = \frac{\gamma r_i N_0 W (1 + Z_l)}{g_{ib(i)}} \qquad (2)$$

Hence, it is customary to limit the *rise over thermal* (ROT), in order to guarantee that the variation in instantaneous transmitted power for each user is small [1]. In other words, limiting ROT provides an alternative condition sufficient (but not equivalent) to Condition C2. Such condition is expressed as

AC1. $\sum_{j=1}^{N} P_j g_{jl} \leq K(W N_0)$,

where K is a fixed value (which can be calculated from ϵ and $P_i^{\max}$, $i = 1, 2, \ldots, N$).

We use condition AC2 to provide a definition of ROT-controlled feasible rates. Similar to above, we reduce the dependency on power vector by introducing the following theorem.

Theorem 2: A vector of rates $\underline{\alpha} = (\alpha_1, \ldots, \alpha_N)$ is an *ROT-controlled feasible rate assignment* iff the vector of effective rates $\underline{\mathbf{r}}$ satisfies

AC1'. $(I - \gamma \mathbf{G}\mathbf{R})^{-1} . \underline{\mathbf{1}}_L \leq (1 + K) \underline{\mathbf{1}}_L$

We denote by Δ_a the *ROT-controlled feasibility region*, i.e. the set of all rate vectors $\underline{\alpha}$ which satisfy condition AC1'.

Remark: Note that if $K \leq \epsilon P_i^{\max}(s_i^{min} + \gamma)$ for $\forall i$, then $\Delta_a \subset \Delta$.

In the remaining sections of this paper, we restrict our attention to Δ_a and ROT-controlled feasibility, since they provide simpler power and rate control mechanisms.

IV. Rate assignment: Maximum Overall Throughput vs. Proportional Fair

In this section, we attempt to introduce the notion of rate assignment. As it is shown above, one can identify a feasibility region, consisting of a (possibly infinite) set of rate vectors, which can be served in a cellular structure. Similar to the last section, we assume gain and base-station assignments are known and fixed. Our goal here is to underline the desirable properties of various rate assignments and to specify the trade-off between fairness and overall throughput. Using such properties, we show that various objective functions can be constructed. Ultimately, we seek to optimize an appropriate objective function over the feasibility region. In other words, after specifying the desirable objective function $F(\underline{\alpha}) := F(\alpha_1, \dots, \alpha_N)$ we seek to solve the following

$$(\alpha_1^{opt}, \dots, \alpha_N^{opt}) := \arg \max_{\underline{\alpha} \in \Delta_a} F(\underline{\alpha}) \tag{3}$$

Note that as discussed before we restrict our attention to the set of ROT-controlled feasible rate vectors, Δ_a. Furthermore, we assume that the objective function is of a social welfare form (see [2]), i.e. it is of the form $\sum_i^N U_i(\alpha_i)$.

Equivalently, the problem can be formulated as an optimization problem with respect to the vector of effective rates as follows.

$$\max \sum_{i=1}^{N} U_i\left(\frac{W r_i}{1-\gamma r_i}\right)$$

subject to $\quad (I - \gamma \mathbf{G}\mathbf{R})^{-1} \underline{\mathbf{1}}_L \leq (1+K)\underline{\mathbf{1}}_L$

A. Overall Throughput

The overall throughput of a system with rate vector $\underline{\alpha}$ is defined to be the sum of rates assigned to mobiles, i.e.

$$U_i^1(\alpha_i) := \begin{cases} \alpha_i & \text{if } \alpha_i < \alpha_i^{max} \\ \alpha_i^{max} & \text{if } \alpha_i \geq \alpha_i^{max} \end{cases} \tag{4}$$

where $w_i \leq W$ is user i's maximum transmission rate.

B. Fairness

The maximum overall throughput might be achieved only at the cost of specific users. In other words, maximizing overall throughput and fairness are to be traded off in most scenarios. To address the fairness, other performance measures such as proportional fairness can be used. The objective function which results in a proportional fair rate assignment is the product of the rates assigned to mobiles [6]. This is equivalent to the sum of the logarithm of rates assigned to mobiles, i.e.

$$U_i^2(\alpha_i) := \begin{cases} \log \alpha_i & \text{if } \alpha_i < \alpha_i^{max} \\ \log \alpha_i^{max} & \text{if } \alpha_i \geq \alpha_i^{max} \end{cases} \tag{5}$$

Note that the proportianl fair, as seen later, is an intermediate solution between an absolute fair (equal rates for all users [3]) and an max-total-throughput assignment.

C. Linearized constraints

The solutions to the above problems require computationally expensive matrix operations and coordination among base station. Our goal in this section is to propose sufficient conditions (perhaps resulting in sub-optimal solutions) with linear structure which allow for a decentralized solution. Theorem 3 provides such conditions.

Theorem 3: A vector of rates $\underline{\alpha} = (\alpha_1, \dots, \alpha_N)$ is an ROT-controlled feasible rate assignment if the vector of effective rates $\underline{\mathbf{r}}$ satisfies

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

Now Theorem 3 can be used to compute a slightly sub-optimal solution with a reduced computational complexity; i.e.

$$\max_{\underline{r} \geq 0} \sum_{i=1}^{N} U\left(\frac{W r_i}{1-\gamma r_i}\right) \tag{6}$$

subject to

$$\sum_{i=1}^{N} \frac{g_{il}}{g_{ib(i)}} r_i \leq \frac{K}{\gamma(1+K)}, \qquad \forall l$$

where function U can be U_1 and U_2 or other utility functions with desirable properties (see [5]).

Consider the Lagrangian

$$\begin{aligned} &\mathcal{L}(\underline{r}, \underline{\mu}) \\ &= \sum_{i=1}^{N} U(r_i) - \sum_{l=1}^{L} \mu_l \left(\sum_{i=1}^{N} \frac{g_{il}}{g_{ib(i)}} r_i - \frac{K}{\gamma(1+K)}\right) \\ &= \sum_{i=1}^{N} \left(U(r_i) - r_i \left(\sum_{l=1}^{L} \frac{g_{il}}{g_{ib(i)}} \mu_l\right) \right) \\ &\quad + \frac{K}{\gamma(1+K)} \sum_{l=1}^{L} \mu_l \end{aligned}$$

If U is monotone non-decreasing and concave function of r (e.g. in case of U_2), the above optimization problem can be addressed by introducing the dual problem. The dual problem is

$$\min_{\underline{\mu} \geq 0} \sum_{i=1}^{N} \phi_i(p_i) + \frac{K}{\gamma(1+K)} \sum_{l=1}^{L} \mu_l \tag{7}$$

where $p_i = \sum_{l=1}^{L} \frac{g_{il}}{g_{ib(i)}} \mu_l$ and

$$\phi_i(p_i) = \max_{r_i} (U(r_i) - r_i p_i). \tag{8}$$

V. Distributed Rate Assignment

Equations (7) and (8) are powerful tools in proposing distributed algorithms in the studied system. In other words, convex duality implies that at the optimum $\underline{\mu^*}$ (which may not be unique), the optimum x^* (maximizing the individual utility minus the cost in (8)) is exactly the solution to the primal problem. Therefore provided the equilibrium prices μ^* can be made to align with the Lagrange multipliers, the individual optima, computed in a decentralized fashion by sources, will align with global optima of (6). A natural way to interpret the above Lagrangian multipliers is to introduce prices for a unit of rate. Note that these prices are not dollar value prices but rather regulating/coordinating signals which are produced by each base station to indicate the level of interference at each sector. In this way, each mobile uses the indication of high level of interference to back off its rate. In summary, under the convexity assumption, the equilibrium points of mobile flow control protocols can be interpreted in terms of mobiles maximizing individual profit based on their own utility functions (U). While base algorithms generate prices to align, exactly or approximately, these "selfish" strategies with social "welfare."

In other words, each mobile varies its rate according to the following expression to maximize its own "profit".

$$\phi_i(p_i) = \max_{r_i}(U(r_i) - r_i \sum_{l=1}^{L} \frac{g_{il}}{g_{ib(i)}}\mu_l). \tag{9}$$

Note that for mobile i, the terms $\frac{g_{il}}{g_{ib(i)}}\mu_l$ are negligible for sectors l who are not in its immediate neighborhood (due to the dominant effect of path loss).

Mobile Algorithm: Each mobile has to compute the above (selfishly) optimal rate at any computation epoch. In order to so, each mobile needs to compute its weighted price (proportional to the sum of its contribution to the ROT at each sector). In CDMA each base station transmits a pilot signal (PS) (with a fixed transmission power P_t^P) over the forward link channel. If forward and reverse link can be assumed to be symmetric (reasonably common assumption), this pilot signal can then be used by mobiles to perform channel estimation and power control. Similarly, we propose that a Pricing Pilot Signal (PPS) is implemented as follows. Each PPS is transmitted on the forward link channel. The transmitted power of PPS for base l is μ_l times the transmission power of the primary pilot signal (PS) P_t^P. Assuming synchronized PPS transmissions from all bases, negligible thermal noise, and also symmetry between forward and reverse links, for each mobile i we have

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

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, while $E_T^P(b(i)) \propto P_t^P g_{ib(i)}$ is the PS energy received by mobile i from its tracking sector $b(i)$. Furthermore, it is required that each mobile knows its own rate utility function $U(.)$. Under these assumptions, the above mobile algorithm can be implemented at each mobile autonomously, using only the available information.

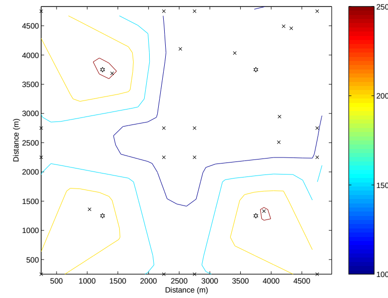

Fig. 1. Rate Assignment Contour Plot, Section VI

Now the challenge is to compute the multipliers (optimal prices) in a distributed manner at each base station over a large network. In other words, we have to provide a decentralized algorithm at each base station for computing the optimal prices.

The simplest algorithm that guarantees these equilibrium prices μ^* is based on the gradient projection :

$$\begin{aligned}\dot{\mu} &= \begin{cases} \beta(\sum_{i=1}^{N} \frac{g_{il}}{g_{ib(i)}} r_i - \frac{K}{\gamma(1+K)}) & \text{if } \mu_l(t) > 0 \\ \beta[\sum_{i=1}^{N} \frac{g_{il}}{g_{ib(i)}} r_i - \frac{K}{\gamma(1+K)}]^+ & \text{if } \mu_l(t) = 0 \end{cases} \\ &= \begin{cases} \frac{\beta}{\gamma}(\sum_{k=1}^{L} \frac{Y_{kl}}{1+Z_k} - \frac{K}{\gamma(1+K)}) & \text{if } \mu_l(t) > 0 \\ \frac{\beta}{\gamma}[\sum_{k=1}^{L} \frac{Y_{kl}}{1+Z_k} - \frac{K}{\gamma(1+K)}]^+ & \text{if } \mu_l(t) = 0 \end{cases}\end{aligned}$$

where the second equality is due to (2) $Y_{kl} = \sum_{i \in M_k} g_{il} P_i$ is the interference receive from sector k at the base station l. Note that if k and l are far from each other $Y_{kl} \approx 0$.

Base Algorithm 1: based on the above, we propose the following algorithm.

$$\dot{\mu} = \begin{cases} \frac{\beta}{\gamma}(\sum_{i=1}^{L} \frac{\hat{Y_{kl}}}{1+\hat{Z_k}} - \frac{K}{\gamma(1+K)}) & \text{if } \mu_l(t) > 0 \\ \frac{\beta}{\gamma}[\sum_{k=1}^{L} \frac{\hat{Y_{kl}}}{1+\hat{Z_k}} - \frac{K}{\gamma(1+K)}]^+ & \text{if } \mu_l(t) = 0 \end{cases}$$

where $\hat{Y_{kl}}$ and $\hat{Z_k}$ are estimates of Y_{kl} and Z_k by base l. In other words, the gradient of the Lagrangian depends on each base station's estimate of the load of its neighboring cells. The better these estimates are, the closer the solution is to the optimal rate assignment.

Base Algorithm 2: Our second algorithm uses a gradient projection method, but uses Condition AC1 instead, in which the prices are constructed in order to stabilize the ROT, i.e. Z_l are matched to threshold K. This algorithm will allow for the same interpretation.

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

This method does not require any estimate of the loading and the interference caused by the neighboring cells.

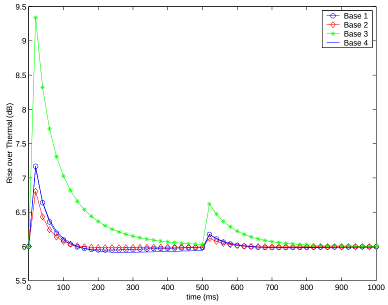

Fig. 2. Rise Over Thermal at Bases for the example in Section VI

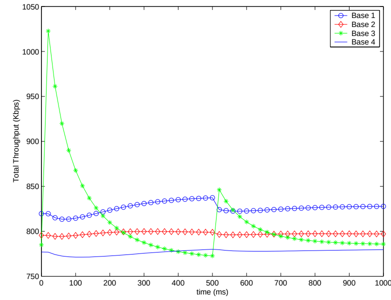

Fig. 3. Total sector throughput for the network in Sections VI

We conjecture that the above two algorithms are equivalent when used with the mobile algorithm. Although we have no proof for this at this moment, we base the conjecture on our observation of various numerical examples in the next section.

VI. Numerical Examples and Simulations

In this section, we provide numerical examples to illustrate the above result. In these examples, we consider the combination of four neighboring cells with asymmetric (non-uniform) mobile locations. Due to limited space, we focus only on one simple instance, ignoring the mobility of the mobiles (in other words, we are assuming that speed of mobiles are less that the channel estimate updates). We assume that base stations are 2500(m) apart, and they update the channel state information every 50 ms. Also each base station l implements Base Algorithm 2 every 20 ms, then uses a pricing pilot signal to broadcast the value μ_l. Mobiles implement the Mobile Algorithm described in Section V. And each mobile is assumed to use the forward link PS and PPS to optimize its utility minus price.

Furthermore, we use a cost-231 propagation model at 1.9 GHz between each mobile and the base stations, in which the exponent for path loss is 3.5 and it includes log-normal shadowing with $\sigma = 8$db. The utility function at each mobile is assumed to be a $\log(\alpha_i)$, where α_i is the mobile's transmission rate. Furthermore, it is assumed that $\gamma = 4dB$, $K = 6dB$, and the total available bandwidth is 1.2 MHz.

Figure 1. shows the contour plot illustrating the rate allocations for various mobiles. To illustrate the dependency of rate on the sum of the gain ratios for neighboring cells we ignored the log normal shadowing when plotting this contour plot. We note that this shows that a proportional fair rate assignment might require unbalanced and unequal rate assignments. This is different from the present implementation in CDMA2000 standards (see [3], [4], [1]). This result shows that an equal rate assignment necessarily will perform sub-optimally with respect to the sector capacity. This is due to the fact that the transmission rate of mobiles at the edge of cells needs to be limited due to their contribution to ROT of multiple cells; but such limitation is unnecessary for mobiles in the inner part of a cell. Figures 2 and 3 show the total throughput and ROT at each base station and in a time varying, dynamic network as described above. Notice that at time $t = 20ms$ a new user arrives and at time $t = 520$ a sudden departure occurs. The figures illustrate the transient behavior of the proposed mobile and base algorithms.

VII. Conclusion

In this paper we address the issue of optimal rate allocation in a wideband CDMA system serving users with varying rates. We first formulated the problem as a global optimization problem, the solution of which depends on the mobile layouts, tracking rules, and each mobile's QoS rate utility. Furthermore, we showed how such a numerical technique for such an optimization problem can provide pairs of distributed mobile/base algorithms whose equilibrium coincide with the solution to the original global optimization problem. In other words, we propose these algorithms as desirable solutions to the issue of decentralized and distributed rate assignment in CDMA systems in scenarios where the changes in the characteristics of the network (layout, tracking assignment, and utilities) are slow. We demonstrate the performance of these algorithms also in realistic scenarios via numerical examples.

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] 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.

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