

Dynamic Rate Control Algorithms for HDR Throughput Optimization

Sem Borst, Phil Whiting
Bell Laboratories, Lucent Technologies
P.O. Box 636, Murray Hill, NJ 07974

Abstract—The relative delay tolerance of data applications, together with the bursty traffic characteristics, opens up the possibility for scheduling transmissions so as to optimize throughput. A particularly attractive approach, in fading environments, is to exploit the variations in the channel conditions, and transmit to the user with the currently ‘best’ channel. We show that the ‘best’ user may be identified as the maximum-rate user when the feasible rates are weighed with some appropriately determined coefficients. Interpreting the coefficients as shadow prices, or reward values, the optimal strategy may thus be viewed as a revenue-based policy.

Calculating the optimal revenue vector directly is a formidable task, requiring detailed information on the channel statistics. Instead, we present adaptive algorithms for determining the optimal revenue vector on-line in an iterative fashion, without the need for explicit knowledge of the channel behavior. Starting from an arbitrary initial vector, the algorithms iteratively adjust the reward values to compensate for observed deviations from the target throughput ratios. The algorithms are validated through extensive numerical experiments. Besides verifying long-run convergence, we also examine the transient performance, in particular the rate of convergence to the optimal revenue vector. The results show that the target throughput ratios are tightly maintained, and that the algorithms are well able to track changes in the channel conditions or throughput targets.

Keywords—Dynamic rate control, fading channels, Quality-of-Service, target throughput ratios, throughput optimization, varying channel conditions.

I. INTRODUCTION

Next-generation wireless networks are expected to support a wide range of services, including high-rate data applications. In contrast to voice users, data applications can usually sustain some amount of packet delay, as long as the throughput over somewhat longer intervals is sufficient. The relative delay tolerance of data applications, together with the bursty traffic characteristics, opens up the potential for scheduling transmissions so as to optimize throughput. A coordinated approach along these lines is proposed in [3].

A related approach may be advocated for low-mobility scenarios such as indoor networks. In such environments, Rayleigh fading frequencies can be quite low, and the fading levels can even be anticipated to some extent. For example, fading can be measured by having the base station provide a pilot signal which can be measured by all the users. These measurements can be fed back to the base station, and used to estimate fading levels and hence user rates in subsequent slots. This is the approach proposed in Qualcomm’s High Data Rate (HDR) scheme [4]. Clearly, it is then advantageous to exploit the variations in the feasible rates, and transmit to the currently ‘best’ user.

With a little simplification, let us suppose that at the start of each slot the base station has perfect knowledge of the maximum feasible rate at which each user can receive and decode a signal with some acceptably low error probability. The question then arises what the ‘best’ user is to be selected for transmission. We show that the ‘best’ user may be identified as the

maximum-rate user when the feasible rates are weighed with some appropriately determined coefficients. Interpreting the coefficients as shadow prices, or reward values, the optimal strategy may thus be viewed as a revenue-based policy. We prove that revenue-based policies optimize throughput relative to pre-specified target ratios. These target values may be set arbitrarily, taking into account the Quality-of-Service requirements of the users, or possibly their current activity levels or locations.

Unfortunately, calculating the optimal revenue vector directly is a complicated problem, requiring detailed information on the channel statistics. Although the feasible rates of the users are assumed known slot by slot, the underlying probability distribution which is producing these rates is unknown. Even if it were known, it would not be easy to use, since the feasible rates might be dependent, so that the computations would be significantly hampered by the curse of dimensionality. To avoid these obstacles, we develop adaptive algorithms for determining the optimal revenue vector on-line in an iterative fashion, without the need for explicit knowledge of the channel behavior.

The application of these algorithms opens up two important possibilities to improve network performance. The first is that admission control can be applied by using a *probing technique*, an approach proposed in [2]. The second possibility is coordinated operation of the base stations in the network, which allows for load sharing and higher throughput for edge users.

The remainder of the paper is organized as follows. In Section II we present a detailed model description, and introduce a class of revenue-based scheduling strategies. We subsequently prove that revenue-based policies optimize throughput relative to pre-specified target ratios, for discrete rate distributions as well as for continuous rates in Sections III and IV, respectively. In Sections V, VI, and VII, we develop adaptive on-line algorithms for determining the optimal revenue vector in an iterative fashion. In Section VIII we describe some numerical experiments which we performed to examine the convergence properties of the proposed control algorithms. We make some concluding remarks in Section IX.

II. MODEL DESCRIPTION

We consider a base station serving M data users. The base station transmits in slots of some fixed duration. In each slot, the base station transmits to exactly one of the users.

We assume that the feasible rates for the various users vary over time according to some stationary discrete-time stochastic process $\{(r_1(n), \dots, r_M(n)), n = 1, 2, \dots\}$, with $r_m(n)$ representing the feasible rate for user m in the n -th slot. We assume that the base station has perfect knowledge of the maximum fea-

In the previous two sections we concluded that revenue-based policies optimize throughput relative to pre-specified target values. However, calculating the optimal revenue vector directly is a complicated problem, requiring detailed information on the channel statistics in the form of the joint stationary distribution of the feasible rates (r_1, \dots, r_M) . Instead, we develop adaptive scheduling algorithms for determining the optimal revenue vector on-line in an iterative fashion without the need for explicit knowledge of the channel behavior.

In the next two sections we assume that the channel state is governed by some discrete-time irreducible Markov chain with a finite discrete state space \mathcal{S} . When the channel state is $s \in \mathcal{S}$, the feasible rates have some continuous M -dimensional distribution $f_s(\cdot)$ on $\mathcal{R} \subseteq [\min, \max]^M$, $0 < \min < \max < \infty$, with zero probability measure in any set of Lebesgue measure zero. In practice, the feasible rates will typically have to be selected from a limited set of discrete values. However, we may adhere to the above assumptions by simply adding a small random perturbation. By choosing the random perturbation sufficiently small, the true achieved throughputs should be arbitrarily close to the perturbed ones.

Denote by $\mathcal{W} := \{ \mathbf{w} \in \mathbb{R}_+^M : \sum_{m=1}^M w_m = 1 \}$ the set of all price vectors. For any $\mathbf{w} \in \mathcal{W}$, denote by $\Xi_m(\mathbf{w})$ the expected average throughput per slot received by user m under price vector \mathbf{w} in stationarity. Define $\Xi_{ave}(\mathbf{w}) := \frac{1}{M} \sum_{m=1}^M \Xi_m(\mathbf{w})$, $\Xi_{\min}(\mathbf{w}) := \min_{m=1, \dots, M} \Xi_m(\mathbf{w})$, and $\Xi_{\max}(\mathbf{w}) := \max_{m=1, \dots, M} \Xi_m(\mathbf{w})$ as the average, the minimum, and the maximum expected throughput per slot under price vector \mathbf{w} over all users, respectively.

Denote by \mathbf{w}^* the optimal revenue vector, i.e., the price vector which balances the expected throughputs. To facilitate the presentation, we assume that \mathbf{w}^* is unique. The analysis may readily be modified for the case where there is a whole range of optimal price vectors.

VI. TWO USERS

We first focus on the case of two users. In the next section, we consider the situation with an arbitrary number of users.

Before describing the algorithm in detail, we first introduce some useful notation. With minor abuse of notation, we write $w = w_1$, so that $w_2 = 1 - w_1$. Denote $\Delta(\mathbf{w}) := r_1(\mathbf{w}) - r_2(\mathbf{w})$, and define $\Delta_N(\mathbf{w}) := \sum_{n=1}^N \Delta(\mathbf{w})$ as the difference in cumulative throughput between users 1 and 2 after N slots. The absolute difference $|\Delta(\mathbf{w})|$ is referred to as the throughput *gap*. We say that the throughput gap *widens* in the n -th slot if $|\Delta(\mathbf{w})| > \max_{n=1, \dots, N-1} |\Delta(\mathbf{w})|$. User 1 is said to be *leading* if $\Delta(\mathbf{w}) > 0$, and is referred to as *lagging* otherwise, and vice versa for user 2. We say that a *cross-over* occurs in the n -th slot if the leading and lagging users exchange positions, which means that the throughput gap changes sign, i.e., $\Delta(\mathbf{w}) (\Delta(\mathbf{w}) - 1) < 0$.

The algorithm may now be described as follows. In every slot, the user with the maximum price-rate product, at the current

price value, is selected for transmission. Thus, the n -th slot is assigned to user 1 if $r_1(\mathbf{w}) > (1 - r_2(\mathbf{w})) r_2(\mathbf{w})$, and to user 2 otherwise (ties being broken arbitrarily).

To drive the price sequence $(\mathbf{w}^{(n)})$ towards the optimal value \mathbf{w}^* , the price is adjusted over time on the basis of the observed throughput realizations. As long as the throughput gap does *not* widen, the price is left unaltered. However, if the throughput gap *does* widen, then the price is changed in favor of the deficit user, thus at the expense of the surplus user. The price of the leading user is decreased by $k^{(n)}$, while the price of the lagging user is simultaneously increased by the same amount.

To ensure convergence, a *reset* is triggered at every cross-over. The step size $k^{(n)}$ is then reduced by incrementing (n) , with $\{k_n, n = 1, 2, \dots\}$ a pre-determined convergent sequence (e.g. $k_{n+1} = k_n^\alpha$ with $\alpha < 1$, or $k_n = k_1^{-\beta}$ with $\beta > 1$).

We now proceed to demonstrate convergence of the above-described algorithm. We first state an important assumption.

Assumption VI.1: (Large-Deviations Assumption)

Let $N_m(\mathbf{w})$ be a random variable representing the average throughput per slot obtained by user m over a period of N slots under price vector \mathbf{w} , given that the initial state of the Markov chain is s . Given a price vector $\mathbf{w} \in \mathcal{W}$ and $\epsilon > 0$, there exist numbers $\xi_m(\mathbf{w}), \xi_m(\mathbf{w}) > 0$ such that for any initial state

$$\mathbb{P}\{ | \frac{N_m(\mathbf{w})}{N} - \Xi_m(\mathbf{w}) | > \epsilon \} \leq \frac{\xi_m(\mathbf{w})}{\epsilon} e^{-D_m^\epsilon(\mathbf{w})N},$$

$$m = 1, 2.$$

It may be verified that the above assumption is satisfied for the feasible-rate process described earlier.

The above assumption ensures almost-sure convergence to the optimal revenue vector as established in the next theorem.

Theorem VI.1: For the scheduling algorithm described above, the price sequence $(\mathbf{w}^{(n)})$ converges to the optimal price \mathbf{w}^* wp 1, and consequently the sequence $(\Delta(\mathbf{w}^{(n)}))$ converges to the optimal value π^{w^*} wp 1.

In preparation for the proof of the above theorem we first present two lemmas.

Lemma VI.1: The price sequence $(\mathbf{w}^{(n)})$ cannot get permanently trapped in either of the intervals $[0, w^* - \epsilon]$ or $[w^* + \epsilon, 1]$.

Proof

We only prove the statement for the interval $[w^* + \epsilon, 1]$. The statement for the interval $[0, w^* - \epsilon]$ follows from symmetry considerations.

The idea of the proof is as follows. As long as the price remains in favor of user 1, the throughput difference continues to have a positive drift, and will wander off to infinity. As a result, the price will keep decreasing in fixed steps, and will eventually turn negative, which is not possible.

For the formal proof details we refer to [5]. \square

Lemma VI.2: The price sequence $(\mathbf{w}^{(n)})$ cannot move from the interval $[0, w^* + \epsilon]$ to the interval $[w^* + 2\epsilon, 1]$ infinitely often. Similarly, $(\mathbf{w}^{(n)})$ cannot move from the interval $[w^* - \epsilon, 1]$ to the interval $[0, w^* - 2\epsilon]$ infinitely often.

Proof

We only prove the first statement, The second one follows from symmetry considerations.

The idea of the proof is as follows. In order for the price sequence to move from the interval $[0, p^* + \epsilon]$ to the interval $[p^* + 2\epsilon, 1]$, it must cross the interval $[p^* + \epsilon, p^* + 2\epsilon]$ from left to right. For that to happen, the algorithm must make a number of ϵ -wrong moves. By an ϵ -wrong move, we mean that the price is increased while the current price is at least ϵ above the optimal value p^* . As will be shown below, the expected number of ϵ -wrong moves before a cross-over occurs is finite. However, as cross-overs occur, the step size will get smaller and smaller, and the required number of ϵ -wrong moves for the interval to be crossed will get larger and larger. As a result, it will eventually become increasingly unlikely for the interval to be crossed.

For the formal proof details we refer to [5]. \square

Proof of Theorem VI.1

Lemma VI.1 implies that the sequence (p_n) spends infinitely many times in the interval $[p^* - \epsilon, 1]$ w.p. 1. Lemma VI.2 shows that the sequence (p_n) returns only finitely many times from the interval $[p^* - \epsilon, 1]$ to the interval $[0, p^* - 2\epsilon]$ w.p. 1. Combining these two statements, we find that the sequence (p_n) spends only finitely many times in the interval $[0, p^* - 2\epsilon]$ w.p. 1. Similarly, we have that the sequence (p_n) spends only finitely many times in the interval $[p^* + 2\epsilon, 1]$ w.p. 1. Hence, for any $\delta > 0$, the sequence (p_n) will eventually enter the interval $[p^* - 2\epsilon, p^* + 2\epsilon]$ w.p. 1, to never leave it again. Thus, the sequence (p_n) converges to the optimal price p^* w.p. 1.

By continuity, the sequence $\mathbb{E}[m(p_n)]$ converges to $\Xi_m(p^*)$, $m = 1, 2$. The convergence of (p_n) then immediately follows. \square

Remark VI.1: Some interesting related algorithms are proposed in [1], [7], [8], [9], [10], where queue lengths instead of rewards are used as weight factors. These algorithms provide throughput guarantees in terms of bounded expected queue lengths (if achievable) rather than target ratios. \square

VII. ARBITRARY NUMBER OF USERS

We now turn to the situation with an arbitrary number of users. In principle, the algorithm described in the previous section for the case of two users may be extended to several users, although there are some subtleties involved in identifying a proper rule for when to trigger a reset.

Here we consider a related but somewhat different algorithm, which may be described as follows. The algorithm makes price updates based on sample periods of pre-determined ever increasing size. Thus, the price updates occur at pre-determined slots (n) , instead of randomly determined slots as before, with $(n) := (n+1) - (n)$ the length of the n -th sample period. In every slot of the n -th sample period, the price vector (p_n) is used for selecting a user for transmission. (From now on we use n to index sample periods, rather than transmission slots as before.)

To drive the price sequence (p_n) towards the optimal point p^* , the price is adjusted over time on the basis of the observed throughput realizations. The *direction* in which the price vector is modified at the n -th update is determined by a random vector (d_n) based on the throughput obtained during the n -th sample period when the price vector (p_n) is used. The *size* of the n -th update is $(\epsilon_n) = \epsilon_{k(n)}$, with $\{k_n, n = 1, 2, \dots\}$ a

pre-determined convergent sequence. Thus, at the $(n+1)$ -th update, the price vector is recursively determined as

$$(n+1) = (n) - (d_n) (\epsilon_n).$$

To ensure convergence, the step size (ϵ_n) is reduced by incrementing (n) every time a reset is triggered. Intuitively, resets should occur rarely far away from the optimal point p^* , but occur readily once the price vector is close to p^* .

It remains to specify the exact rules for (i) how to determine the update direction (d_n) , and (ii) when to trigger a reset.

(i) For every user the empirical average throughput over the sample period is computed. The users are then partitioned into two groups: (a) those with above-average throughput; (b) those with below-average throughput. The prices of the above-average users are decreased, while the prices of the below-average users are increased. As the sample size grows, so that with high probability the empirical average throughputs line up with the true expected throughputs, this ensures that the price vector gets closer to the optimal point p^* in some appropriate sense, as will be shown later.

Formally, the procedure may be described as follows. Denote by m_n the throughput received by user n during a particular sample period in which price vector (p_n) is used. Define $\bar{m}_{ave} := \frac{1}{M} \sum_{m=1}^M m_n$ as the average throughput over all users. Denote by $\Omega^- := \{n : m_n \leq \bar{m}_{ave}\}$ and $\Omega^+ := \{n : m_n > \bar{m}_{ave}\}$ the groups of below-average and strictly above-average users, respectively. Then the price update direction (d_n) is determined as

$$i(d_n) = \frac{i}{\sum_{m \in \Omega^-} m} \in \Omega^-, \quad (6)$$

$$j(d_n) = \frac{-j}{\sum_{m \in \Omega^+} m} \in \Omega^+. \quad (7)$$

Note that the price ratios within both Ω^- and Ω^+ are maintained. This ensures that the expected throughput of the below-average users increases, while the expected throughput of the above-average users decreases, as may be easily checked. In case $\Omega^+ = \emptyset$, the price vector is simply left unaltered.

Also note that the above price update cannot be applied in case the price values of some of the users in Ω^+ are zero. To prevent that situation from happening, the price process will be restricted to the set $\mathcal{W}_\nu := \{p \in \mathcal{W} : p_n \geq \nu \text{ for all } n = 1, \dots, N\}$, with $\nu := \min(\nu_{\min} / (\nu_{\min} + (\nu_{\max} - 1)\nu_{\max}))$. It is easily verified that $p^* \in \mathcal{W}_\nu$. In order to restrict the price process to the set \mathcal{W}_ν , the update is truncated at the boundary if necessary.

(ii) To ensure convergence, a reset is triggered under the condition that every user has been a member of Ω^+ at least once during a consecutive sequence of updates. Once the reset has occurred, the next one is not triggered until every user has been a member of Ω^+ at least once again.

We now proceed to prove convergence of the above-described algorithm. We first discuss a few important assumptions.

Large-deviations assumption

As described above, the algorithm works by making price updates based on samples of ever increasing size. To ensure

convergence, we need that as the sample size grows, a ‘correct’ price update direction is selected with sufficiently high probability. Given a price vector $\mathbf{p} \in \mathcal{W}$, user i is called ϵ -below-average (respectively, ϵ -above-average) if $\Xi_m(\mathbf{p}) < \Xi_{ave}(\mathbf{p}) - \epsilon$ (respectively, $\Xi_m(\mathbf{p}) > \Xi_{ave}(\mathbf{p}) + \epsilon$). We say that the price update direction is ‘ ϵ -right’ if all the ϵ -below-average users belong to Ω^- and have their price increased, and all the ϵ -above-average users belong to Ω^+ and have their price decreased. This ensures that the price vector gets closer to the optimal point \mathbf{p}^* in some appropriate sense, as will be shown later. Now remember that at each update, the prices of the empirical below-average users are increased, while the prices of the empirical above-average users are decreased. Thus, for the price update direction to be ‘correct’, it is critical that the empirical average throughputs line up with the true expected throughputs. This then motivates the following assumption.

Assumption VII.1: (Large-Deviations Assumption)

Let $\bar{r}_m(\mathbf{p})$ be a random variable representing the average throughput per slot obtained by user m over a period of $L(\mathbf{p})$ slots under price vector \mathbf{p} in stationarity. Given a price vector $\mathbf{p} \in \mathcal{W}$ and $\epsilon > 0$, there exist a ϵ -neighborhood $\xi(\mathbf{p})$ of \mathbf{p} and numbers $\xi_m(\mathbf{p})$, $\xi_m(\mathbf{p}) > 0$ such that

$$\mathbb{P}\{|\bar{r}_m(\mathbf{p}') - \Xi_m(\mathbf{p})| > \xi_m(\mathbf{p})\} \leq \xi_m(\mathbf{p}) e^{-D_m^\xi(\mathbf{p})L(\mathbf{p})},$$

for all $\mathbf{p}' \in \xi(\mathbf{p})$, $m = 1, \dots, M$.

We refer to [5] for a proof that the above assumption is satisfied for the feasible-rate process described earlier.

Boundary conditions

We further require that when a correct price direction is selected, the update cannot be truncated to an arbitrarily small size. The following assumption implies that if a correct price direction is chosen, then for small enough step size δ , the price sequence will stay away from the boundary.

Assumption VII.2: There exist positive constants $\epsilon^* > 0$, $\delta^* > 0$ such that for all price vectors $\mathbf{p} \in \mathcal{W}_\nu$, for any ϵ -right direction (\cdot) , and for any $\delta \in (0, \delta^*)$,

$$\mathbf{p} + \delta(\cdot) \in \mathcal{W}_\nu.$$

To check that the above assumption is satisfied, it suffices to verify that extremely low prices cannot be decreased and that extremely high prices cannot be increased. First consider a user with a price $p_i < \min / (\min + (\epsilon - 1) \max)$. Then the throughput of user i is zero, which means that the price of user i is increased if the price direction is right. Similarly, the throughput of a user with a price $p_j > \max / (\min + \max)$ is ϵ -above-average for some $\epsilon > 0$, so that the price of user j is decreased if the price direction is right.

Function $\Phi(\cdot)$

As indicated above, we also need that when a correct price update direction is selected, the price vector gets closer to the optimal point \mathbf{p}^* by some definite amount. To measure distance from \mathbf{p}^* , we introduce a function $\Phi(\cdot)$ which attains a unique minimum at \mathbf{p}^* . Define $\Gamma_\epsilon := \{\mathbf{p} \in \mathcal{W} : \Xi_{\max}(\mathbf{p}) - \Xi_{\min}(\mathbf{p}) \leq \epsilon\}$ as an ‘ ϵ -neighborhood’ of \mathbf{p}^* . The following assumption implies that if a correct price update direction is chosen, then outside Γ_ϵ the reduction in the value of $\Phi(\cdot)$ for small

enough step size δ , is at least ϵ times some constant of proportionality ϵ .

Assumption VII.3: There exist positive constants $\epsilon^* > 0$, $\delta^* > 0$, $\epsilon > 0$ such that for all price vectors $\mathbf{p} \notin \Gamma_\epsilon$, for any ϵ -right direction (\cdot) , and for any $\delta \in (0, \delta^*)$,

$$\Phi(\mathbf{p} + \delta(\cdot)) \leq \Phi(\mathbf{p}) - \epsilon.$$

We will consider two alternative choices for the function $\Phi(\cdot)$. The first one is

$$\Phi(\mathbf{p}) := \Xi_{\max}(\mathbf{p}) - \Xi_{\min}(\mathbf{p}),$$

i.e., the maximum difference in expected throughput between any pair of users. By definition $\Phi(\mathbf{p}^*) = 0$, and $\Phi(\mathbf{p}) \geq 0$ for all $\mathbf{p} \neq \mathbf{p}^*$, with strict inequality in case the optimal price vector \mathbf{p}^* is unique.

The second function that we will consider is

$$\Phi(\mathbf{p}) := \sum_{m=1}^M m \Xi_m(\mathbf{p}),$$

i.e., the total expected revenue earned. As found in Section III, the optimal price vector \mathbf{p}^* minimizes that quantity over all vectors in the set \mathcal{W} , i.e., $\Phi(\mathbf{p}^*) \leq \Phi(\mathbf{p})$ for all $\mathbf{p} \in \mathcal{W}$, $\mathbf{p} \neq \mathbf{p}^*$, with strict inequality in case \mathbf{p}^* is unique.

We refer to [5] for a proof that Assumption VII.3 is indeed satisfied for the above two $\Phi(\cdot)$ functions.

In contrast to the first one, the second $\Phi(\cdot)$ function is also suitable to show that Assumption VII.3 is satisfied for various alternative options to select a price update direction, for example

$$i^* = 1 - \epsilon > 0 \quad \mathbf{p}^* = \arg \min_{m=1, \dots, M} m, \quad (8)$$

$$j^* = -1 - \epsilon < 0 \quad \mathbf{p}^* = \arg \max_{m=1, \dots, M} m, \quad (9)$$

and $k = n / (\epsilon - 2)$ for all $n \neq n^*$, n^* for n a given positive sequence with $\lim_{n \rightarrow \infty} n = 0$. In the sequel this will be referred to as the ‘Update-Extreme’ algorithm, as opposed to the procedure described earlier which will be called the ‘Move-to-Average’ algorithm.

The next theorem establishes almost-sure convergence to the optimal revenue vector \mathbf{p}^* for the Move-to-Average algorithm.

Theorem VII.1: The price sequence $\mathbf{p}(n)$ converges to the optimal price vector \mathbf{p}^* wp 1, and consequently the sequence $\Phi(\mathbf{p}(n))$ converges to the optimal value $\Phi(\mathbf{p}^*)$ wp 1.

The above theorem is proved in [5]. The proof for the Update-Extreme algorithm is mostly similar, except for a somewhat different notion of a correct price update direction.

Remark VII.1: In the present paper we focus on establishing almost-sure convergence to the optimal revenue vector \mathbf{p}^* . This critically relies on the step sizes $\{\delta_k, k = 1, 2, \dots\}$ being a convergent sequence. As an alternative, the step sizes may be kept fixed at some given value δ . We expect that the price sequence will then continue to oscillate around \mathbf{p}^* , but with smaller amplitudes for smaller values of δ . Observe however that there is an inherent trade-off between the accuracy achieved on the one hand and the speed the convergence, and thus the responsiveness to changing conditions, on the other hand. The value of δ then may be used to find the right balance between these two conflicting objectives. \square




Fig. 1. Normalized expected throughput $\Xi_i(w)$ as function of w .

VIII. NUMERICAL RESULTS

In this section we describe some numerical experiments which we conducted to investigate the convergence properties of the proposed control algorithms. Besides verifying long-run convergence, we also examine the transient performance, in particular the rate at which the prices converge to the optimal values.

In the first three experiments we consider continuous rate distributions. In the fourth experiment we assume a discrete distribution where the feasible rates are determined by a fading process via the signal-to-noise ratio. The fading process is modeled using a discrete number of sinusoidal oscillators as described by Jakes' model [6].

In the final three experiments, we examine how well the throughput ratios are maintained, and how well the algorithms are able to track changes in the channel conditions or throughput targets.

A. Two users with exponential rates

In the first experiment we consider a model of two users with independent rates. The feasible rate for user i is governed by a conditional exponential distribution on some interval $[r_{\min}, r_{\max}]$, i.e.,

$$f_i(r) = \gamma_i^{-1} [1 - e^{-\gamma_i(r - R_{\min})}], \quad r \in [r_{\min}, r_{\max}],$$

with $\gamma_i = 1 - e^{-\gamma_i(R_{\max} - R_{\min})}$ a normalization coefficient, $\gamma_i = 1, 2$. We take $[r_{\min}, r_{\max}] = [10, 400]$ Kbits/s and assume $(\gamma_1, \gamma_2) = (0.02, 0.01)$. Thus, the feasible rate for user 2 is about twice as large in distribution as for user 1. The throughput target for user 2 is also set twice as large as for user 1, i.e., $(\gamma_1, \gamma_2) = (1, 2)$.

The normalized expected throughputs for these parameters as a function of w are plotted in Figure 1. From the figure, we see that the optimal price is $w^* \approx 0.6$, which may more precisely be determined as $w^* \approx 0.593$ using bisection.

We ran the control algorithm described in Section VI for 1000 slots. We used step sizes $w_{k+1} = w_k - \beta$, with initial value $w_1 = 0.5$ and reduction factor $\beta = 0.9$. The resulting price trajectories




Fig. 2. Price trajectory for 2 users over 1000 slots.




Fig. 3. Price trajectories for 2 users vs. w^* (non-geometric step sizes).

are graphed in Figure 2 for a period of 1000 slots. Observe that the prices converge to the optimal values in roughly 300 slots, which corresponds to about 0.3 seconds of operation.

We repeated the above experiment for non-geometric step sizes $\beta_k = \beta^{-k}$, with β successively chosen as 1.5, 2.0, 3.0, 4.0. Note that the sum of the price changes is still convergent, although the step sizes decay slower than before. The corresponding price trajectories are shown in Figure 3 for a period of 1000 slots. We see that convergence is considerably slower for smaller values of β , i.e., slower decay of the step sizes.

B. Three users

In the second experiment we consider a scenario with three users. As before, the feasible rate for user i follows a conditional exponential distribution on the interval $[10, 400]$ with parameters $(\gamma_1, \gamma_2, \gamma_3) = (0.02, 0.01, 0.02)$. Thus, the feasible rate for user 2 is about twice as large in distribution as for users 1 and 3.

The target throughput ratios for the three users are set equal, i.e., $(\gamma_1, \gamma_2, \gamma_3) = (1, 1, 1)$. The optimal revenue vector is $w^* \approx (0.424, 0.152, 0.424)$ as may be determined using numerical integration and two-dimensional bisection. Observe that the




Fig. 4. Price trajectories for 3 users over 5000 slots vs. w^* (Move-to-Average algorithm).




Fig. 6. Price trajectories for 8 users with 15,000 slots vs. w^* (Move-to-Average algorithm).




Fig. 5. Price trajectories for 3 users over 5000 slots vs. w^* (Update-Extreme algorithm).

optimal price for users 1 and 3 is higher than for user 2, as is required in order to obtain equal throughput, since the feasible rate for user 2 is stochastically larger.

We ran the two control algorithms described in Section VII for 5000 slots, or approximately 5 seconds of operation, with $(k) = 10$ slots for the k -th update. This amounts to roughly 30 price updates. The initial revenue vector is set to $(1) = (0.3, 0.6, 0.1)$. We used step sizes $k = -2$, $k = 1, 2, \dots$. The resulting price trajectories are depicted as the solid curves in Figures 4 and 5. The revenue vector for the Update-Extreme algorithm after 30 price updates is $(30) \approx (0.441, 0.123, 0.436)$, quite close to the optimal one.

We repeated the above experiment for the Update-Extreme algorithm using 40 and 60 slots for the k -th update, with the same power series for k . The corresponding price trajectories are reproduced as the dashed lines in Figure 5 for user 1 in the first case and user 2 in the second (with similar results for the remaining prices.) As expected, we see that using fewer samples per price update leads to a slower and ‘noisier’ convergence to

TABLE I
FEASIBLE RATE PER SLOT AS FUNCTION OF SNR.

Signal-to-Noise Ratio (dB)	Rate (bits)
$-5.0 < \text{SNR}$	1000
$-10.0 < \text{SNR} \leq -5.0$	500
$-20.0 < \text{SNR} \leq -10.0$	250
$-30.0 < \text{SNR} \leq -20.0$	100
$\text{SNR} \leq -30.0$	30

the optimal revenue vector w^* .

C. Eight users

In the third experiment we consider a situation with eight users. As before, the feasible rate for user k follows a conditional exponential distribution on the interval $[10, 400]$. The exponents were chosen at random uniformly in $[0.01, 0.05]$, and turned out to be approximately $(0.0489, 0.0263, 0.0139, 0.0480, 0.0220, 0.0107, 0.0461, 0.0128)$.

The target throughput ratios are again set equal for all users. As before, we expect that a larger value of the exponent α_k , inducing smaller feasible rates, requires a higher price in order to obtain equal throughput.

We ran the two control algorithms described in Section VII for 15,000 slots, or approximately 15 seconds of operation, with $(k) = 30$ slots for the k -th update. This amounts to roughly 55 price updates. The initial revenue vector is set at random. We used step sizes $k = -2$, $k = 1, 2, \dots$. The resulting price trajectory for the Move-to-Average algorithm is graphed in Figure 6.

D. Discrete rates driven by a fading process

We now consider a case with discrete rates governed by independent fading processes as described by Jakes’ model [6]. The mean received powers of user 1, 2 and 3 are -15.0 dB, 0.0 dB, and -10.0 dB, respectively. The feasible rates per slot then follow from Table 1.




Fig. 7. Empirical average throughput for 3 users over 10,000 slots (Move-to-Average algorithm with $\delta_k = k^{-3/2}$).

The throughput target for user 2 is set twice as large as for users 1 and 3, i.e., $(\gamma_1, \gamma_2, \gamma_3) = (1, 2, 1)$.

We ran the two control algorithms described in Section VII for 10,000 slots, with k slots for the k -th update. We used step sizes $\delta_k = k^{-3/2}$ and $\delta_k = k^{-2}$, $k = 1, 2, \dots$

As explained earlier, the discrete rate values are perturbed by adding a small uniformly distributed random variable to obtain a continuous version of the problem. We thus ensure that the optimal control algorithm is determined by the revenue vector only.

The empirical average throughputs are depicted in Figures 7, and 8. The achieved throughputs under the Update-Extreme algorithm are approximately 130 bits per slot for both users 1 and 3 and 270 bits per slot for user 2, quite close to the target ratios. Under the Move-to-Average algorithm the realized throughputs are reasonably close to the target ratios too, provided the step size is reduced sufficiently slowly as in Figure 7.

The corresponding price trajectories are displayed in Figures 9, and 10. We see that that under the Update-Extreme algorithm the prices converge to the optimal values in about 5 seconds. Under the Move-to-Average algorithm the prices converge fairly quickly too, unless the step size is reduced so fast that the process gets essentially overdamped.

E. Comparison with a forcing scheme

We now compare the revenue-based algorithms with a forcing scheme. The forcing scheme assigns the k -th transmission slot to the user $i^*(k)$ with the current minimum normalized throughput, i.e.,

$$i^*(k) = \arg \min_{m=1, \dots, M} m(k) / \gamma_m.$$

By construction, the forcing scheme realizes the target throughput ratios perfectly, in the sense that with probability 1,

$$\frac{i(k)}{j(k)} \rightarrow \frac{\gamma_i}{\gamma_j}, \quad \text{as } k \rightarrow \infty$$

for all pairs of users $i, j = 1, \dots, M$.




Fig. 8. Empirical average throughput for 3 users over 10,000 slots (Update-Extreme algorithm with $\delta_k = k^{-2}$).




Fig. 9. Price trajectories for 3 users over 10,000 slots (Move-to-Average algorithm with $\delta_k = k^{-3/2}$).




Fig. 10. Price trajectories for 3 users over 10,000 slots (Update-Extreme algorithm with $\delta_k = k^{-2}$).




Fig. 11. Empirical average throughput for 3 users over 5000 slots (forcing algorithm).

The downside of the forcing scheme of course, is that it generally achieves lower throughput in absolute terms, as it does not take advantage of the variations in the feasible rates. Under i.i.d. assumptions, the throughput obtained under the forcing scheme may in fact be computed in closed form as

$$i(\cdot) \rightarrow i^*, \quad \text{as } \rho \rightarrow \infty.$$

$$\text{with } i^* = \sum_{j=1}^M \lambda_j / \mathbb{E}[\lambda_j].$$

We repeated the experiment of the previous subsection for the forcing scheme. The empirical average throughputs are reproduced in Figure 11 for a period of 5000 slots. The achieved throughputs are approximately 90 bits per slot for both users 1 and 3, and 180 bits per slot for user 2. The results show how tightly the target throughput ratios are maintained under the forcing scheme. In absolute terms however, the throughput for all users is about 30% smaller than for the revenue-based algorithms.

F. Tracking capability

We now examine how well the algorithms are able to track sudden changes in the target throughput ratios or channel conditions. In the first experiment, the throughput target for user 3 is initially set to some low value. After 80 seconds, the throughput target is suddenly incremented to allow for the transmission of a data burst for user 3.

The resulting price trajectory for the Move-to-Average algorithm is plotted in Figure 12. The optimal price values for the new throughput ratios are also indicated as dashed straight lines. The results show that after a few oscillations the prices quickly settle down to the new optimal values.

IX. CONCLUSION

We considered the problem of scheduling data users with varying channel conditions so as to obtain the optimal long-run throughputs for given target ratios. We have shown that the problem may be solved by selecting users for transmission according to an optimal revenue vector λ^* which balances the




Fig. 12. Price adjustment to allow for data burst for user 3 (Move-to-Average algorithm).

expected throughputs. We presented a wide class of stochastic control algorithms which ensure almost-sure convergence to λ^* and thus achieve the optimal long-run throughputs.

Numerical experiments showed that the convergence to the optimal revenue vector is in practice quite rapid (of the order of a few seconds), and that the algorithms have the ability to track changes in the channel conditions and throughput targets. Further experiments are required to determine which form of the algorithm is most adequate for implementation in the HDR scheme. The algorithms may also be enhanced by allowing the step sizes or the sample sizes to be adapted in response to non-stationary changes in the feasible rate declarations

REFERENCES

- [1] Andrews, D.M., Kumaran, K., Ramanan, K., Stolyar, A.L., Vijayakumar, R., Whiting, P.A. (2000). CDMA data QoS scheduling on the forward link with variable channel conditions. Technical Memorandum, Bell Laboratories, Lucent Technologies.
- [2] Bambos, N., Chen, S.C., Mitra, D. (1995). Channel probing for distributed access control in wireless communication networks. In: *Proc. IEEE Globecom '95*.
- [3] Bedekar, A., Borst, S.C., Ramanan, K., Whiting, P.A., Yeh, M. (1999). Downlink scheduling in CDMA data networks. In: *Proc. IEEE Globecom '99*, 2653–2657.
- [4] Bender, P., Black, P., Grob, M., Padovani, R., Sindhusayana, N., Viterbi, A. (2000). CDMA/HDR: a bandwidth-efficient high-speed wireless data service for nomadic users. *IEEE Commun. Mag.* **38**, 70–77.
- [5] Borst, S.C., Whiting, P.A. (2000). Dynamic rate control algorithms for HDR throughput optimization. Technical Memorandum 10009626-000828-21TM, Bell Laboratories, Lucent Technologies. Submitted for publication.
- [6] Jakes, W.C. (1974). Multipath interference. In: *Microwave Mobile Communications*. Ed. W.C. Jakes, IEEE Press, Piscataway.
- [7] Kahale, N., Wright, P.E. (1997). Dynamic global packet routing in wireless networks. In: *Proc. IEEE Infocom '97*, 1441–1421.
- [8] Shakkottai, S., Stolyar, A.L. (2000). A study of scheduling algorithms for a mixture of real and non-real time data in HDR. Technical Memorandum, Bell Laboratories, Lucent Technologies.
- [9] Shakkottai, S., Stolyar, A.L. (2000). Throughput-optimal scheduling for time-varying channels: the exponential rule. Technical Memorandum, Bell Laboratories, Lucent Technologies.
- [10] Tassiulas, L. (1998). Linear complexity algorithms for maximum throughput in radio networks and input queued switches. In: *Proc. IEEE Infocom '98*, 553–559.