Title
A Novel Coordinated Connection Access Control and Resource Allocation Framework for 4G Wireless Networks

Journal
IEEE/ACM Transactions on Networking, 23(4)

ISSN
1063-6692 1558-2566

Authors
Yang, Chao
Jordan, Scott

Publication Date
2015-08-01

DOI
10.1109/TNET.2014.2326168

Peer reviewed

A Novel Coordinated Connection Access Control and Resource Allocation Framework for 4G Wireless Networks

Chao Yang and Scott Jordan, *Member, IEEE*

***Abstract*—In the academic literature on cellular network design, resource allocation algorithms often attempt to maximize total utility or throughput over a short time period, and connection access control often admits arrivals if and only if there are sufficient resources. In this paper, we investigate how connection access control and resource allocation can be coordinated to jointly achieve maximum total utility. We propose a decomposition in which resource allocation maximizes long-term average utility for each system state, and connection access control maximizes long-term average utility over all system states. We discuss the resulting interface and give examples of algorithms that satisfy this decomposition and interface. Simulation results illustrate that the optimal connection access control policy may block applications with relatively low average utility per unit rate even when capacity is available, and that coordinated connection access control and resource allocation can outperform uncoordinated approaches.**

***Index Terms*—Connection access control, coordination, resource allocation.**

I. Introduction

HERETOFORE, connection access control (CAC) and resource allocation (RA) have been designed to accomplish different goals in cellular networks. Resource allocation algorithms typically attempt to maximize the total utility of all active users, e.g., the total number of voice users or the total throughput of data users. In contrast, connection access control algorithms typically admit a new connection if and only if it is believed that capacity is available to ensure acceptable performance. Thus, while RA focuses on utility, CAC typically ignores utility and merely focuses on capacity.

Here, we investigate how CAC and RA can be coordinated to both focus on utility. There is ample reason to believe that coordination of CAC and RA could be beneficial. First, the volume of data traffic on 4G networks has surpassed that of voice, and video is quickly becoming the dominant traffic class by volume. It is expected that voice, data, and video will all be important revenue generators. Whereas the first three generations of cellular networks segregated capacity for voice and data applications, future networks will transmit all applications over the Internet Protocol (IP) and thereby share capacity among all application classes. Multiplexing of traffic classes with different quality-of-service (QoS) requirements presents a challenge to CAC and RA. When capacity was segregated, CAC and RA for voice could both focus on maximizing the number of voice users, and CAC and RA for data could both focus on maximizing throughput. In contrast, when capacity is shared, it no longer is meaningful for CAC to admit a new connection if and only if it is believed that capacity is available to ensure acceptable performance since such an admit decision for a voice call may result in an unacceptably high marginal decrease in utility for active data and video connections. Instead, we suggest that CAC should also focus on maximizing utility.

Second, orthogonal frequency-division multiplexing (OFDM) is used in 4G networks. The total bandwidth is divided into a set of orthogonal narrowband subcarriers. Different users' channels on a particular subcarrier are often uncorrelated due to independent fading. Thus, OFDM systems often adaptively allocate subcarriers among multiple users based on instantaneous channel information [1]. The benefit of this approach increases with the number of users per cell. However, if too many users are admitted, the performance of each user decreases. Thus, coordinated CAC and RA may be particularly beneficial in OFDM systems.

There is a great deal of research literature on uncoordinated resource allocation and connection access control for cellular networks. Resource allocation algorithms usually attempt to maximize total throughput (see, e.g., [2]–[5]) or total utility (see, e.g., [6]–[9]). In some cases, the allocation attempts to satisfy the QoS requirements of the application, e.g., a minimum rate requirement for voice applications. Connection access control has typically depended upon the application type. For voice, CAC usually admits a new call if and only if there are available resources (see, e.g., [10]–[12]). In contrast, data applications are often treated as not having any QoS requirement, and thus not requiring any CAC. CAC for video applications has been less addressed and remains an open problem.

However, there is little research literature on coordinated resource allocation and connection access control for cellular networks. A number of papers propose a weak type of coordination in which CAC admits voice users subject to available capacity and RA attempts to maximize the total utility of data users subject to performance constraints on voice users, see, e.g., [13].

Manuscript received May 14, 2013; revised February 06, 2014; accepted May 04, 2014; approved by IEEE/ACM TRANSACTIONS ON NETWORKING S. Puthenpura. Date of publication June 02, 2014; date of current version August 14, 2015.
The authors are with University of California, Irvine, Irvine, CA 92697 USA (e-mail: Chao.Yang@uci.edu; sjordan@uci.edu).
Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.
Digital Object Identifier 10.1109/TNET.2014.2326168

Other papers propose coordination by treating the goal as one of maximizing sum throughput of voice and data subject to QoS requirements, see, e.g., [14].

The existing RA and CAC algorithms have several obvious drawbacks. First, RA principally focuses on instantaneous or short-term throughput or utility. These algorithms typically do not attempt to maximize *long-term* performance. Second, CAC algorithms often simply admit users if there is sufficient capacity (see, e.g., [10]–[14]). However, the utility seen by each user is often a function of the number of active users, and admitting an arrival may cause a decrease in the total system utility. In addition, in systems with multiple classes of users, admitting an arrival of one class may result in a forced blocking in the near future of a user of a different class that may have resulted in higher utility or revenue. Third, the goals of current CAC and RA algorithms are different. The goal of CAC is often to evaluate if there are enough resources to admit arrivals. In contrast, the goal of RA is usually to maximize instantaneous or short-term throughput or total utility. Cooperation between CAC and RA is limited, and the system may not achieve *long-term* optimal performance.

Here, we propose coordinated CAC and RA on the basis of user utility. First, we suggest a joint optimization of *long-term* average user utility over both CAC and RA policies. We define the utility as the function of average rate within certain time window. This guarantees that CAC and RA have the same goal. Due to the high complexity of an exhaustive search for the optimal solution, we then propose a decomposition into separate CAC and RA problems, which make the cooperation between CAC and RA simple and efficient. Whereas traditional RA in the literature typically attempts to maximize total utility over a relatively short time period (see, e.g., [6]–[9]), we propose that RA should attempt to maximize the *long-term* average utility *for each system state*, where the state is defined as number of active applications of each class. Whereas traditional CAC in the literature typically admits new connections if and only if there are available resources, we propose that CAC should attempt to maximize the long-term average utility *over all system states*.

We discuss the resulting interface between RA and CAC. The RA module evaluates the feasible region and the average total utility of each state and passes this information to the CAC module. Based on this information, the CAC module selects the admission policy, and based on this policy, decides whether to admit an arrival of each user class. Then, the CAC module passes the admission decision to the RA module, completing the feedback loop.

We also give examples of RA and CAC algorithms that satisfy this decomposition and interface. We show how stochastic dynamic programming can be used to find the optimal CAC policy. The optimal admission decision thus takes into account potential future utility. Using numerical examples, we illustrate that the optimal policy may not be to admit an arrival if there are sufficient resources and may block applications with relatively low average utility per unit rate even when capacity is available.

The rest of this paper is as follows. In Section II, we define a user's channel, rate, and utility. In Section III, we formulate a joint CAC and RA problem and propose a decomposition into separate CAC and RA problems. In Section IV, we provide an implementation example to explain how to evaluate the average utility of each state and how to use stochastic dynamic programming to design an connection access control policy. Finally, in Section V, the performance of our framework is illustrated by numerical simulation results.

II. System Model

We consider a single-cell downlink OFDM system with N subcarriers. The bandwidth B of each subcarrier is assumed to be less than the coherence bandwidth of the channel so that the channel response can be considered flat. The rate of user k on subcarrier n at time t is

$$r_{k,n,t}(p_{k,n,t}) = B\log_2\left(1 + p_{k,n,t}\frac{|H_{k,n,t}|^2}{\sigma^2 + I}\right)$$

where $p_{k,n,t}$ is the power allocated, $|H_{k,n,t}|^2$ is the composite channel gain, σ^2 is the noise power, and I is the interference power. The channel gain $|H_{k,n,t}|^2 = \alpha_{k,n,t}^2\gamma_{k,t}PL_{k,t}$ is composed of fast fading $\alpha_{k,n,t}^2$ that changes significantly in sequential time periods, slow fading and shadowing $\gamma_{k,t}$ whichthat changes little in sequential time periods but may change significantly during a few seconds, and path loss $PL_{k,t}$ that depends on user position and changes significantly during tens of seconds.[1] Fast fading on different subcarriers is assumed to be independent to each other. The total rate of user k at time t is the sum of the user's rate over all subcarriers

$$R_{k,t} = \sum_{n=1}^{N} r_{k,n,t}.$$

Assume there are L classes of applications. Each application class is associated with a utility function U_l. We propose that users measure the satisfaction of applications by the average rate achieved within a time window W_l, rather than by the instantaneous rate. The size of W_l is determined by the application—e.g., for video applications, W_l should correspond to several *group-of-pictures*, and for voice applications, W_l should be tens of milliseconds (the minimum time scale over which humans can perceive differences in voice quality).

Denote the application of user k by $b_k = l$, and user k's average rate at time t within the most recent window W_{b_k} by $S_{k,t} = \int_{\tau=t-W_{b_k}}^{t} R_{k,\tau}/W_{b_k}d\tau$. User utility is thus a function of average rate $U_{b_k}(S_{k,t})$. We wish to consider elastic applications (e.g., data) modeled by concave utility, semi-elastic applications (e.g., video) modeled by sigmoid utility, and inelastic applications (e.g., voice) modeled by step utility; examples of each are shown in Fig. 1.

The QoS requirement of an application is modeled by a minimum rate S'_{b_k}. Elastic applications do not have any rate requirement, so $S'_{b_k} = 0$. The minimum rate of an inelastic application corresponds to the location of the step. For semi-elastic applications, the utility generated per unit rate is maximized at tan-

[1]Here, we assume perfect instantaneous channel estimation, and that the modulation and coding achieve the capacity of the channel. Consideration of the effects of imperfect channel estimation, delayed channel estimation, and the particular modulation and coding scheme is important, but is outside the scope of this paper. In particular, analysis of adaptive modulation and coding (see, e.g., [15]) in conjunction with our approach would be interesting.

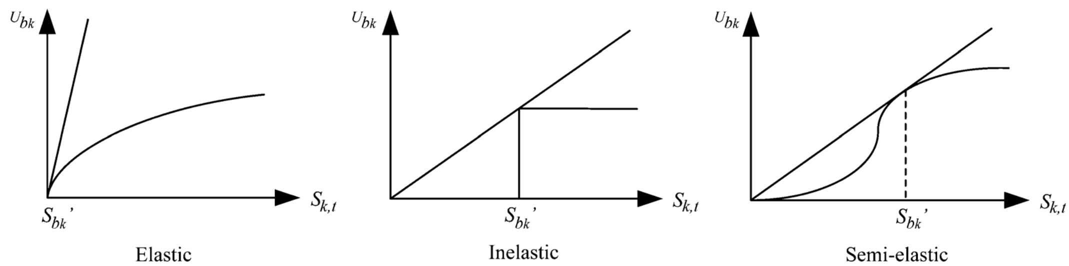

Fig. 1. Utility function.

TABLE I
NOTATION

Notation	Description
N	number of subcarriers
B	subcarrier bandwidth
P	downlink power of base station
t	current time
ι	discretized variable notation
Δt_s	time slot length
t^ι	time slot number, $t^\iota = \lceil t/\Delta t_s \rceil$
M	number of slices in the partition of combined pathloss and shadowing
L	number of application classes
W_l	length of time window of class l application
b_k	the application class of user k
U_{b_k}	utility function of user k
$S_{b_k}^{'}$	minimum rate requirement of user k
$\gamma_{k,t}$,γ_{k,t^ι}	shadowing of user k
$PL_{k,t}$,PL_{k,t^ι}	pathloss of user k
$\psi_{k,t}$, ψ_{k,t^ι}, ψ_k, $\psi_k^{(m)}$	combined pathloss and shadowing of user k
$\alpha_{k,n,t}^2$, $\alpha_{k,n}^2$	fast fading of user k
$\|H_{k,n,t}\|^2$, $\|H_{k,n,t^\iota}\|^2$	composite channel fading of user k
$p_{k,n,t}$, p_{k,n,t^ι}, $p_{k,n}^{(m)}$	allocated power to user k on subcarrier n
I and σ^2	interference power and noise power
$r_{k,n,t}$	rate of user k on subcarrier n
$R_{k,t}$, R_{k,t^ι}	rate of user k at time t, at slot t^ι
$S_{k,t}$, S_{k,t^ι}, S_k, $S_k^{(m)}$	average rate of user k within preceding time window
μ_t^ι	power price
$\overline{\lambda}_{k,t^\iota}$, $\lambda_k^{(m)}$	average rate price
β_l and η_l	intensity of arrival and departure process.
$\mathbf{x}$	The state of the network
$\mathbf{g}$	CAC policy
G, $G_\mathbf{x}$	CAC policy space
Q^{CAC}, Q^{RA}	CAC and RA policy
Λ, Λ^ι	feasible region

gent point $S_{b_k}' = \arg\max_{S_{b_k}} U_{b_k}/S_{b_k}$. Thus, it is reasonable to assume that a compression algorithm would be designed on the assumption that the average rate is maintained above this threshold.[2]

Arrivals of applications of class l are assumed form independent Poisson processes with intensity β_l, and the duration of connections of class l are assumed to be independent and exponentially distributed with mean $1/\eta_l$. Our notation is summarized in Table I.

III. Joint Framework

In this section, we first formulate a joint CAC and RA problem. Then, we propose a decomposition into separate CAC and RA problems and discuss the resulting interface between CAC and RA.

A. Joint CAC and RA Optimization Problem

We begin by consideration of the policy spaces for CAC and RA. For CAC, we assume that departures are never blocked, so that the policy space can be written as

$$G = \{\mathbf{g} = (g_1, g_2, \ldots, g_L) : g_l \in \{0,1\},\ l = 1, \ldots, L\}$$

where $g_l = 0$ (resp. $g_l = 1$) denotes that an arrival of class l should be blocked (resp. admitted). We call a user *active* if the user has been admitted to the system and has not yet departed. Denote the time of the jth event (arrival or departure) by $t^{(j)}$, and the set of active users at time t by

$$D_t = \{k | \text{user } k \text{ is active at time } t\}.$$

For each arrival event j, the CAC policy is given by $\mathbf{g}^{(j)} = \{g_1^{(j)}, g_2^{(j)}, \ldots, g_L^{(j)}\}$. The state of the network is represented by a vector $\mathbf{x} = \{x_1, x_2, \ldots, x_L\}$ where x_l is the number of active class-l applications; the state just after event j is denoted by $\mathbf{x}^{(j)} = \{x_1^{(j)}, x_2^{(j)}, \ldots, x_L^{(j)}\}$. The CAC policy, denoted by

$$Q^{\mathrm{CAC}} = \mathbf{g}^{(j)}\left(\mathbf{x}^{(j-1)}\right)$$

thus decides on admission of a connection on the basis of the state immediately before the connection's arrival. The RA policy, denoted by

$$Q^{\mathrm{RA}} = \left\{p_{k,n,t} \big| |H_{k,n,\tau}|^2,\ \forall k, n,\ \tau \le t \text{ and } \mathbf{x}^{(j)}\right\}$$

assigns powers to each user and subcarrier at each time t on the basis of each user's historical channels and on the current state.

We now turn to the desired optimization metric. It is common, in both multiclass connection access control problems and resource allocation problems, to consider efficiency and/or fairness. In connection access control, efficiency is often defined in terms of a weighted sum of admitted connections of the traffic classes, or equivalently as a weighted sum of blocking probabilities of traffic classes (see, e.g., [16]–[18]). In contrast, fairness is often defined either as admitting all connections subject to performance constraints or as equalizing blocking probabilities of the traffic classes (see, e.g., [19]–[21]). Efficiency and fairness are often viewed as a tradeoff, and hence some approaches try to balance the two goals (see, e.g., [22]).

Similarly, in resource allocation, efficiency is often defined in terms of a weighted sum of throughput of throughput of the traffic classes (see, e.g., [3] and [23]), whereas fairness is often

[2]Alternatively, one might assume that the inflection point is a reasonable threshold.

defined as equalizing throughput of connections or of traffic classes (see, e.g., [24] and [25]). Again, these are often viewed as a tradeoff, and some approaches try to balance them (see, e.g., [26] and [27]). Some papers adopt the concept of utility, which can be a function of any performance metric including throughput. One advantage of utility is that it can be a nonlinear function of the performance metric, and can thus represent more complex goals than maximization of weighted throughput.

Resource allocation algorithms intended that use utility typically attempt to maximize the total utility of all active users during a short time period (see, e.g., [6]–[9]), and connection access control algorithms typically admit a new connection if and only if it is believed that capacity is available to ensure acceptable performance (see, e.g., [10]–[14]). Such uncoordinated RA and CAC has several disadvantages. First, if RA focuses on short-term performance, then long-term utility is not necessarily maximized. Second, in traditional CAC algorithms, the goal is to admit as many users as possible if there is enough capacity. However, this algorithm is suboptimal since the expected utility from the new arrival may well be less than the decrease in utility of current users resulting from their decrease in average future rate. Moreover, when residual capacity is low and a relatively low-paying application class arrives, admitting the new arrival may cause the system to lose the chance to admit a future arrival of a higher-paying user. Finally, the goals of RA and CAC are different, and the lack of cooperation limits long-term performance. We thus propose that both RA and CAC should focus on long-term user utility. By focusing on utility, CAC and RA have the same goal. When residual capacity is low and a relatively low-paying application class arrives, CAC can cooperate with RA to judge whether it is optimal to block the arrival with the hope that a higher-paying application may arrive soon.

There are several options for this metric. Whereas traditional RA in the literature typically attempts to maximize total utility over a relatively short time period, we propose that RA should attempt to maximize the *long-term* average utility *for each system state*. Define the duration of state $\mathbf{x}^{(j)}$ as $\Delta t^{(j)} = t^{(j+1)} - t^{(j)}$. Denote the average utility per unit time of state $\mathbf{x}^{(j)}$ by

$$U_{\mathrm{avg}}\left(\mathbf{x}^{(j)}, \Delta t^{(j)}\right) = \frac{1}{\Delta t^{(j)}} \int_{t^{(j)}}^{t^{(j+1)}} \sum_{k \in D_t} U_{b_k}(S_{k,t})dt.$$

Note that this metric thus averages utility over different time periods in which the state is the same, as opposed to maximizing short-term throughput.

Whereas traditional CAC in the literature typically admits new connections if and only if there are available resources, we propose that CAC should attempt to maximize the long-term average utility *over all system states*. The joint connection access control and resource allocation problem is thus

$$\max_{Q^{\mathrm{CAC}}, Q^{\mathrm{RA}}} \lim_{J \uparrow \infty} \frac{1}{t^{(J+1)}} \sum_{j=0}^{J} U_{\mathrm{avg}}\left(\mathbf{x}^{(j)}, \Delta t^{(j)}\right) \Delta t^{(j)}$$

where J denotes the total number of events.

The same optimization metric is thus used by both CAC and RA. When a new user arrives, CAC does not make an admission decision based only on current available capacity; it will cooperate with RA to evaluate both current and long-term utility. It is worth noting that the policy of admitting all users and maximizing total throughput is a special case in which $\mathbf{g}^{(j)}(\mathbf{x}^{(j-1)}) = \mathbf{1}\ \forall \mathbf{x}^{(j-1)}$ and $W_l = 1, U_l(S_{k,t}) = R_{k,t}\ \forall l$.

This formulation, however, requires the joint determination of the optimal CAC and RA policies. An exhaustive search among all possible CAC and RA policies within the joint policy space is almost certainly infeasible. In Section III-B, we thus attempt to decompose the problem into separate CAC and RA problems. The challenge is to formulate decomposed problems that retain the desired coordination and that can be implemented using a simple interface between the two modules.

B. Decomposition and Interface

We start with resource allocation. Traditional resource allocation focuses on one period of time and optimizes system performance during this period. We argue that resource allocation should attempt to maximize the long-term average utility for each system state, i.e., there should be a RA policy for each state $\mathbf{x}$. In an infinite length of time period, state $\mathbf{x}$ occurs during time

$$T(\mathbf{x}) = \bigcup_{\left\{j \le J : \mathbf{x}^{(j)} = \mathbf{x}\right\}} \left[t^{(j)}, t^{(j+1)}\right).$$

The average utility of state $\mathbf{x}$ over an infinite time period is thus

$$U_{\mathrm{avg}}(\mathbf{x}) = \lim_{J \uparrow \infty} \frac{1}{|T(\mathbf{x})|} \int_{T(\mathbf{x})} \sum_{k \in D_t} U_{b_k}(S_{k,t})dt$$

where $|\ |$ denotes cardinality.

For each state $\mathbf{x}$, resource allocation should maximize long-term average utility, and the resource allocation problem can be written as

$$\begin{aligned} \overline{U}_{\mathrm{avg}}(\mathbf{x}) &= \max_{Q^{\mathrm{RA}}} U_{\mathrm{avg}}(\mathbf{x}) \\ \text{s.t.}\quad & \sum_{k \in D_t} \sum_{n=1}^{N} p_{k,n,t} \le P \qquad \forall t \in T(\mathbf{x}) \\ & p_{k,n,t} \ge 0 \qquad \forall k \in D_t, n,\ t \in T(\mathbf{x}) \\ & S_{k,t} > S'_{b_k} \qquad \forall k \in D_t,\ t \in T(\mathbf{x}) \end{aligned} \tag{1}$$

where P is the base station's downlink transmission power.

We define the feasible region Λ as the set of states for which there exists a resource allocation that satisfies the constraints in (1)

$$\Lambda = \left\{\mathbf{x} : \exists \mathrm{Q}^{\mathrm{RA}} \text{ which can solve problem (1)}\right\}.$$

We turn next to connection access control. While RA attempts to maximize the long-term average utility for each state, CAC should attempt to maximize the long-term average utility over all states. The major question is the following: Does CAC need detailed knowledge of the RA algorithm, or can it treat RA as a black box? The literature does not provide a clear answer to this question. To strengthen the modularity of the decomposition, we propose that CAC should treat RA as a black box and exchange a minimal set of information. The question remains as to what information will suffice. Does CAC need to know the power and

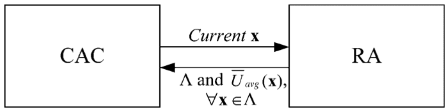

Fig. 2. Joint CAC and RA flowchart.

subcarrier allocation? The resulting QoS? The utility earned as a function of time? We propose that CAC only needs to know Λ, the feasible region, and $\{\overline{U}_{\mathrm{avg}}(\mathbf{x})\ \forall \mathbf{x} \in \Lambda\}$, the set of average utilities earned in each state by the resource allocation policy. Does RA need to know the history of arrivals and departures? We propose that RA only needs to know the current state. The resulting interface is illustrated in Fig. 2.

Based on this interface, the decomposed CAC problem becomes one of choosing an admission policy that maximizes long-term average utility over the entire feasible region

$$\max_{Q^{\mathrm{CAC}}} \lim_{J\uparrow\infty} \frac{1}{t^{(J+1)}} \sum_{j=0}^{J} \overline{U}_{\mathrm{avg}}\left(\mathbf{x}^{(j)}\right) \Delta t^{(j)}. \tag{2}$$

In summary, the RA module determines the feasible region Λ and estimates the maximum average utility of each state $\overline{U}_{\mathrm{avg}}(\mathbf{x}^{(j)})$, either by estimating utility as a function of the channel distribution or by direct observation of the real system. With this information, the CAC module can determine the optimal admission policy. Both of these roles can be done offline. Then online, when a user arrives, the CAC module determines whether to admit merely using a lookup based on state, and the RA module determines power and subcarrier allocation based on the current state and the current set of channel gains.

Cooperation is thus instilled between CAC and RA by sharing the same optimization metric of long-term average utility. The CAC and RA modules can treat each other as black boxes and communicate limited information via a simple interface. This decomposition provides flexibility to design CAC and RA policies separately and also reduces complexity.

IV. Implementation

In this section, we give examples of RA and CAC algorithms that satisfy the decomposition and interface proposed in Section III. This implementation will provide methods to estimate $\overline{U}_{\mathrm{avg}}(\mathbf{x})$ and to use stochastic dynamic programming to determine the optimal CAC policy.

A. Resource Allocation

Whereas the model above assumed continuous time, in real systems, resources are allocated during each time-slot [28]. We thus would like to formulate a discrete-time version of the RA problem. Define Δt_s as the duration of a time-slot, $t^\iota = \lceil t/\Delta t_s \rceil$, $t^{(j),\iota} = \lceil t^{(j)}/\Delta t_s \rceil$, and $[t^{(j),\iota},\ t^{(j+1),\iota})$ as $\{t^\iota | t^{(j),\iota} \le t^\iota < t^{(j+1),\iota}\}$. Then, the set of time-slots in which the system is in state $\mathbf{x}$ is

$$T^\iota(\mathbf{x}) = \bigcup_{\{j \le J : \mathbf{x}^{(j)} = \mathbf{x}\}} \left[t^{(j),\iota}, t^{(j+1),\iota}\right)$$

and the discretized average utility is

$$U^\iota_{\mathrm{avg}}(\mathbf{x}) = \lim_{J\uparrow\infty} \frac{1}{|T^\iota(\mathbf{x})|} \sum_{T^\iota(\mathbf{x})} \sum_{k\in D_{t^\iota}} U_{b_k}(S_{k,t^\iota})$$

where $S_{k,t^\iota} = \sum_{\tau^\iota = t^\iota - W_k + 1}^{t^\iota} R_{k,\tau^\iota}/W_k$.

Denote the total number of active users while in state $\mathbf{x}$ by $K(\mathbf{x}) = \sum_{l=1}^{L} x_l$. In different epochs j within $T^\iota(\mathbf{x})$, the total number of users is the same, but the active user set D_{t^ι} may be different. This makes it difficult to calculate $\overline{U}^\iota_{\mathrm{avg}}(\mathbf{x})$. Thus, for each state $\mathbf{x}$, we propose to consider a fixed set of users over an infinite length of time. For simplicity of notation, for each state $\mathbf{x}$, we renumber the users from 1 to $K(\mathbf{x}) = \sum_{l=1}^{L} x_l$ and restart time from 1 to T^ι,[3] and we therefore drop the $\mathbf{x}$ from $K(\mathbf{x})$.

Denote the power allocation by $\mathbf{p} = \{p_{k,n,t^\iota},\ \forall k, n, t^\iota\}$. Each subcarrier can be allocated to at most one user, thus denote the feasible set of power and subcarrier allocations by $\mathbf{A} = \{\mathbf{p}\ \text{s.t.}\ \forall t^\iota, n,\ p_{k,n,t^\iota} > 0$ for at most one user $k\}$. The discretized maximum average utility in state $\mathbf{x}$ is then

$$\begin{aligned} \overline{U}^\iota_{\mathrm{avg}}(\mathbf{x}) = \max_{\mathbf{p}\in\mathbf{A}} & \frac{1}{T^\iota} \sum_{t^\iota=1}^{T^\iota} \sum_{k=1}^{K} U_{b_k}(S_{k,t^\iota}) \\ \text{s.t.} \quad & \sum_{k=1}^{K} \sum_{n=1}^{N} p_{k,n,t^\iota} \le P \qquad \forall t^\iota;\ p_{k,n,t^\iota} \ge 0\ \forall k, n, t^\iota \\ & S_{k,t^\iota} > S'_{b_k} \qquad \forall k, t^\iota. \end{aligned} \tag{3}$$

We solve a dual problem by introducing a set of intermediate variables $\mathbf{d} = \{d_{k,t^\iota},\ \forall k, t^\iota\}$ as bounds on the achieved rates, i.e., $S_{k,t^\iota} \ge d_{k,t^\iota}\ \forall k, t^\iota$. The Lagrange is

$$\begin{aligned} F(\mathbf{d}, \mathbf{p}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = & \frac{1}{T^\iota} \sum_{t^\iota=1}^{T^\iota} \sum_{k=1}^{K} U_{b_k}(d_{k,t^\iota}) \\ & + \sum_{t^\iota=1}^{T^\iota} \sum_{k=1}^{K} \lambda_{k,t^\iota}(S_{k,t^\iota} - d_{k,t^\iota}) + \sum_{t^\iota=1}^{T^\iota} \mu_{t^\iota}\left(P - \sum_{k=1}^{K} \sum_{n=1}^{N} p_{k,n,t^\iota}\right) \end{aligned} \tag{4}$$

where $d_{k,t^\iota} > S'_{b_k}$, $\boldsymbol{\lambda} = \{\lambda_{k,t^\iota},\ \forall k, 1 \le t^\iota \le T^\iota\}$ are the Lagrangian multipliers associated with the rate constraints which are interpreted as rate prices [29], and $\boldsymbol{\mu} = \{\mu^\iota_t, 1 \le t^\iota \le T^\iota\}$ are the Lagrangian multipliers associated with the power constraints.

However, knowing the channel information for all $1 \le t \le T^\iota$ is impractical, and we propose to estimate $\overline{U}^\iota_{\mathrm{avg}}(\mathbf{x})$ using the channel distribution. Define a random variable $\alpha^2_{k,n}$ representing fast fading for user k on subcarrier n, whose distribution is presumed known; denote $\boldsymbol{\alpha^2} = \{\alpha^2_{k,n},\ \forall k, n\}$. Combine slow fading, shadowing, and path loss for user k into a single random variable ψ_k, whose distribution is presumed known; denote $\boldsymbol{\psi} = \{\psi_k,\ \forall k\}$. It will be helpful to think of the power allocations as random variables $p_{k,n}$ and the achieved average rates as random variables S_k that are functions of the random variables $\boldsymbol{\alpha^2}$ and $\boldsymbol{\psi}$ and of the resource allocation policy. Denote the set of shadowing and path losses for all users except user k

[3] We also presume that each user has been in the system since time $t^\iota = -(W_{b_k} - 2)$, so that S_{k,t^ι} is defined at $t^\iota = 1$.

by $\boldsymbol{\psi}_{-k} = \{\psi_{\hat{k}},\ \forall \hat{k} \neq k\}$. We assume that the distribution of combined slow fading and shadowing ψ_k is independent of k. As $T^{\iota} \uparrow \infty$

$$\frac{1}{T^{\iota}} \sum_{k=1}^{K} \sum_{t^{\iota}=1}^{T^{\iota}} U_{b_k}(S_{k,t^{\iota}}) \rightarrow \sum_{k=1}^{K} E_{\boldsymbol{\alpha^2},\boldsymbol{\psi}} U_{b_k}(S_k).$$

Partition the domain of ψ_k into M slices, with the lower bound of slice m denoted by $\psi^{(m)}$. Denote the probability of slice m by $q_m = Pr(\psi_k = \psi^{(m)})$. When in slice m, denote the power allocated to user k on subcarrier n by $p_{k,n}^{(m)}$ and the average rate of user k by $S_k^{(m)}$. The optimal policy $\mathbf{p}$ is then given by set of Lagrangian multipliers for each user denoted by $\{\lambda_k^{(m)},\ \forall k, m\}$. The optimization problem (3) thus becomes

$$\begin{aligned} \overline{U}^{\iota}_{\text{avg}}(\mathbf{x}) \approx & \max_{\left\{\lambda_k^{(m)}, \forall k, m\right\}} \sum_{k=1}^{K} \sum_{m=1}^{M} U_{b_k}\left(E_{\boldsymbol{\alpha^2},\boldsymbol{\psi}_{-k}} S_k^{(m)}\right) q_m \\ \text{s.t.} \quad & \sum_{k=1}^{K} \sum_{m=1}^{M} \sum_{n=1}^{N} E_{\boldsymbol{\alpha^2},\boldsymbol{\psi}_{-k}} \left(p_{k,n}^{(m)}\right) q_m \leq P \\ & E_{\boldsymbol{\alpha^2},\boldsymbol{\psi}_{-k}} S_k^{(m)} > S'_{b_k} \qquad \forall k, m. \end{aligned} \tag{5}$$

The derivation is in Appendix I. The discretized feasible region is

$$\Lambda^{\iota} = \left\{\mathbf{x} : \exists \left\{\lambda_k^{(m)}, \forall k, m\right\} \text{ which can solve problem (5)}\right\}.$$

This optimization problem can be solved using a subgradient algorithm. The resulting complexity is polynomial in the number of users K.

B. Connection Access Control

It remains to formulate and solve a discrete-time version of problem (2). With the discretized feasible region, the CAC policy space becomes

$$G_{\mathbf{x}} = \{\mathbf{g} \in G : g_l = 0 \text{ if } \mathbf{x} + \mathbf{e}_l \notin \Lambda^{\iota}\}$$

where $\mathbf{e}_l$ is a vector of zeros, except for one for the lth component. The CAC policy $Q^{\text{CAC}} = \mathbf{g}^{(j)}(\mathbf{x}^{(j-1)})$ must decide on admission of a connection on the basis of the state immediately before the connection's arrival.

The state $\mathbf{x}^{(j)}$ is a continuous-time Markov chain. It can be converted into a discrete-time Markov chain via standard uniformization techniques [30, pp. 110, 209] as follows. The expected time in state $\mathbf{x}^{(j)}$ under policy $\mathbf{g}$ is

$$\nu_{\mathbf{x}}(\mathbf{g}) = \left\lceil \frac{1}{\sum_{l=1}^{L} \beta_l g_l + \sum_{l=1}^{L} x_l \eta_l} \cdot \frac{1}{\Delta t_s} \right\rceil \Delta t_s \tag{6}$$

and a bound on these expected times is

$$\nu = \left\lceil \frac{1}{\sum_{l=1}^{L} (\beta_l + \overline{x}_l \eta_l)} \cdot \frac{1}{\Delta t_s} \right\rceil \Delta t_s \tag{7}$$

where $\overline{x}_l$ is the maximum number of class-l applications in the network that can be calculated by allocating all the resources to class-l applications. The embedded discrete-time Markov chain transition probabilities are given by

$$P_{\mathbf{x}\mathbf{x}'} = \begin{cases} x_l \eta_l \nu_{\mathbf{x}}(\mathbf{g}),\ l = 1, \ldots, L, & \text{if } \mathbf{x}' = \mathbf{x} - \mathbf{e}_l \\ \beta_l g_l \nu_{\mathbf{x}}(\mathbf{g}),\ l = 1, \ldots, L, & \text{if } \mathbf{x}' = \mathbf{x} + \mathbf{e}_l \\ 0, & \text{else.} \end{cases}$$

An equivalent discrete-time problem is

$$\max_{Q^{\text{CAC}}} \lim_{J \uparrow \infty} \frac{1}{J} \sum_{j=1}^{J} \overline{U}^{\iota}_{\text{avg}}\left(\mathbf{x}^{(j)}\right).$$

The optimal CAC policy can then be determined using stochastic dynamic programming. The corresponding value iteration algorithm [30, p. 210] is

$$\begin{aligned} V_i(\mathbf{x}) = \max_{\mathbf{g} \in G_{\mathbf{x}}} \Bigg\{ & \overline{U}^{\iota}_{\text{avg}}(\mathbf{x}) + \frac{\nu}{\nu_{\mathbf{x}}(\mathbf{g})} \sum_{\mathbf{x}' \in \Lambda^{\iota}} P_{\mathbf{x}\mathbf{x}'}(\mathbf{g}) V_{i-1}(\mathbf{x}') \\ & + \left(1 - \frac{\nu}{\nu_{\mathbf{x}}(\mathbf{g})}\right) V_{i-1}(\mathbf{x}) \Bigg\} \end{aligned} \tag{8}$$

where i denotes the iteration number, $V_i(\mathbf{x})$ denotes the expected return of state $\mathbf{x}$, and $\mathbf{x}'$ is a possible next state.

In summary, the offline portion of the CAC module uses stochastic dynamic programming to determine the optimal connection control policy; this is then used online to decide upon admissions. Upon admitted arrivals or departures, the CAC module passes the new state $\mathbf{x}$ to the RA module. Offline, the RA module uses problem (5) to determine a set of rate prices $\{\lambda_k^{(m)}\ \forall k, m\}$; users in the same slice m using the same application l will face identical rate prices, denoted $\lambda^{(l,m)}$. These rate prices can be used online to allocate power and subcarriers based on the current state and the current set of channel gains using interpolation; if $b_k = l$, then an interpolated rate price can be set to

$$\overline{\lambda}_{k,t^{\iota}} \approx \begin{cases} \lambda^{(l,m)} - \left(\lambda^{(l,m)} - \lambda^{(l,m-1)}\right) & \\ \quad \cdot \frac{\psi_{k,t^{\iota}} - \psi^{(m)}}{\psi^{(m-1)} - \psi^{(m)}}, & \text{if } \psi^m < \psi_{k,t^{\iota}} < \psi^{m-1}, \\ & \text{or } \psi_{k,t^{\iota}} < \psi^M \\ \lambda^{(l,m)} - \left(\lambda^{(l,m+1)} - \lambda^{(l,m)}\right) & \\ \quad \cdot \frac{\psi_{k,t^{\iota}} - \psi^{(m)}}{\psi^{(m)} - \psi^{(m+1)}}, & \text{if } \psi_{k,t^{\iota}} > \psi^1. \end{cases}$$

The complexity of the stochastic dynamic programming problem is not easy to express in terms of the size of the problem. The number of iterations required is discussed in Section V. However, this portion is completed offline, and the online calculation is very simple.

C. Analysis of CAC Policy

In this section, we characterize the near-optimal CAC policy for a special case of two classes of applications. There is a large amount of academic literature on CAC policies in various types of networks. A small proportion of these papers have succeeded at providing a characterization of the CAC policy that maximizes a specified metric among a specified class of policies. Foschini *et al.* [31] considered a system with two classes of applications requiring identical amounts of a shared single resource type. It proved that the policy that maximizes average revenue, among the class of coordinate convex policies, consists of a threshold on one of the application classes. Ross and Tsang [32] considered systems with two classes of applications that may require different amounts of a shared single resource type. It proved under certain conditions that the policy that maximizes average revenue, among the class of coordinate convex policies, may consist of a threshold on one or both of the application classes. Altman *et al.* [33] considers both a more general class of policies (Markov policies) and a more general class of single resource systems (with two or more classes of applications sharing a single type of resource). It proves that sometimes the optimal application classes can be ordered in terms of when it is optimal to admit them.

However, almost all of the literature characterizes optimal policies only in systems with a single type of resource. In the model considered in this paper, there are two types of resources: power and subcarriers. Jordan and Varaiya [34] considered systems with multiple classes of applications that may require different amounts of multiple types of resources. It proved that the policy that maximizes average revenue, among the class of coordinate convex policies, has a certain type of threshold structure.

However, in all of these CAC policy characterizations, the resource requirement is fixed during the duration of a connection, whereas in the model considered in this paper the resource requirement varies during the connection.

It appears to be very difficult to characterize the optimal CAC policy in systems with multiple application classes sharing multiple resources, particularly when resource requirements vary during the connection. We thus consider an approximated version of the optimization problem as follows. Remove the ceiling function from the expected time in state $\mathbf{x}^{(j)}$ in (6) and from its corresponding bound in (7). This approximation allows the maximization in the value iteration (8) to be replaced by a term that depends on whether $V_{i-1}(\mathbf{x}+\mathbf{e}_l) > V_{i-1}(\mathbf{x})$

$$\begin{aligned} V_i(\mathbf{x}) = V_{i-1}(\mathbf{x}) + \overline{U}^{\iota}_{\text{avg}}(\mathbf{x}) \\ + \nu \sum_{l=1}^{L} \{ x_l \eta_l \left[V_{i-1}(\mathbf{x}-\mathbf{e}_l) - V_{i-1}(\mathbf{x}) \right] \\ + \beta_l \mathbf{1} \left(V_{i-1}(\mathbf{x}+\mathbf{e}_l) > V_{i-1}(\mathbf{x}) \right) \\ \cdot \left[V_{i-1}(\mathbf{x}+\mathbf{e}_l) - V_{i-1}(\mathbf{x}) \right] \} \end{aligned} \tag{9}$$

where the terms including $V_{i-1}(\mathbf{x}-\mathbf{e}_l)$ are included only if $x_l > 0$ and the terms including $V_{i-1}(\mathbf{x}+\mathbf{e}_l)$ are included only if $\mathbf{x}+\mathbf{e}_l \in \Lambda^{\iota}$.

This simplification makes it possible to provide a characterization of the optimal CAC policy for a special case of two classes of applications.

Theorem 1: If: 1) there are $L = 2$ classes of applications; 2) the feasible region Λ^{ι} is a rectangle $\{(x_1, x_2)$ such that $0 \le x_1 \le \bar{x}_1, 0 \le x_2 \le \bar{x}_2\}$; and 3) $\overline{U}^{\iota}_{\text{avg}}(\mathbf{x})$ is monotonically increasing with x_1 for all fixed x_2, then the CAC policy that satisfies (9) always admits class-1 arrivals while inside the feasible region.

The proof is given in Appendix II. An immediate corollary of Theorem 1 is that if there are $L = 2$ classes of applications, if the feasible region Λ^{ι} is a rectangle $\{(x_1, x_2)$ such that $0 \le x_1 \le \bar{x}_1, 0 \le x_2 \le \bar{x}_2\}$, and if $\overline{U}^{\iota}_{\text{avg}}(\mathbf{x})$ is monotonically increasing with both x_1 and x_2, then the CAC policy that satisfies (9) admits all arrivals while inside the feasible region.

It is worth noting that the condition that the feasible region Λ^{ι} is a rectangle cannot be ignored. In Section V, we give an example in which the other conditions in the theorem are satisfied but the feasible region is not a rectangle, and show that the optimal policy is not as simply characterized.

V. Simulation Results

In this section, we examine the performance of the proposed decomposition via simulation. We assume there are $N = 1000$ subcarriers, with each subcarrier occupying a bandwidth of $B = 10$ kHz. The base station's downlink transmission power for all applications is 46 dBm. The path loss, including the antenna gain, is [28]

$$PL_k = 128.1 + 37.6 \log 10(\zeta_k) + 21 \log 10 \left(\frac{f_c}{2.0} \right) - 15 \text{ dB}$$

where ζ_k is the distance from the user to the base station and $f_c = 2000$ MHz is the central frequency. The shadowing follows a lognormal distribution with mean value 0 dB and variance 10 dB [28]. The domain of $\psi_k = \gamma_k PL_k$ is partitioned into $M = 35$ slices. Fast fading $\alpha^2_{k,n}$ follows an Exponential distribution with mean value 1. The duration of one time-slot is $\Delta t_s = 1$ ms [35], and fast fading is generated every three time-slots independent of previous fading.

Users move at a constant speed of 10 km/h with direction determined by a random walk [36] through a hexagonal cell with intersite distance equal to 750 m. If users stay in the cell for a long period of time and/or if users' speed is high enough, these users are approximately uniformly distributed in a donut around the base station from $\zeta = 0.01$ km to $\zeta = 0.25$ km.

We consider three types of applications: inelastic (voice), elastic (data), and semi-elastic (video). However, for purposes of illustration, in Sections V-A–V-C, we consider dynamic sharing of only two classes of applications at a time. The total system resources (power and subcarriers) are thus assumed to be statically partitioned between the two classes of focus and the remaining unmodeled classes.

A. Two Semi-Elastic Classes

We first focus on dynamic sharing of system resources between two classes of semi-elastic (video) applications. We presume that 40% of system resources are allocated to these two classes, i.e., $N = 400$ subcarriers and $P = 42$ dbm downlink power. Noise plus interference $\sigma^2 + I = 7.6232 * 10^{-28}$ mW. The time window over which utility is evaluated is $W_1 = W_2 = 133$ slots, corresponding a common choice for the length of one group of pictures in MPEG4. The utility function for a class-l application is

$$U_l(S_{k,t^{\iota}}) = \begin{cases} a_l (S_{k,t^{\iota}})^2, & \text{if } S_{k,t^{\iota}} < S_l^f \\ c_l (S_{k,t} + b_l)^{\frac{1}{3}}, & \text{else} \end{cases}$$

where $S_{k,t^{\iota}}$ is also expressed in units of 100 kb/s, $a_1 = 4/5 * (2/5)^{1/3}/(12/5)^2$, $b_1 = -2$, $c_1 = 4/5$, $S_1^f = 240$ kb/s, $a_2 = (5/6)^{1/3}/25$, $b_2 = -25/6$, $c_2 = 1$, $S_2^f = 500$ kb/s. Both utility functions are sigmoid [29]; we set the minimum rate requirements to the rate at the maximum average utility, which is given by $S_1' = 300$ kb/s and $S_2' = 625$ kb/s. The arrival processes are Poisson processes with intensities $\beta_1 = 1/12/s$ and $\beta_2 = 1/30/s$, respectively, and the connection durations

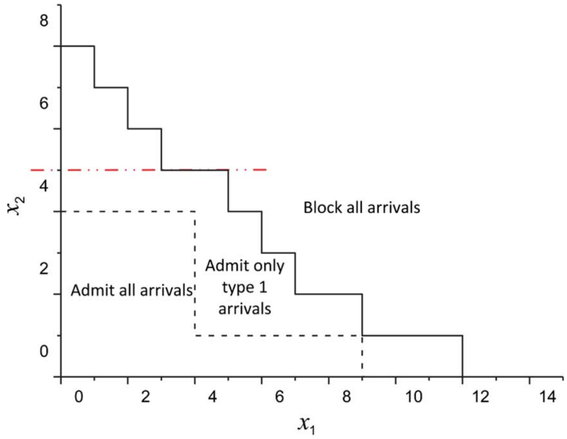

Fig. 3. Feasible region and CAC policy for two semi-elastic classes.

are exponentially distributed with means $1/\eta_1 = 180$ s and $1/\eta_2 = 300$ s, respectively.

We use the value iteration algorithm to determine the connection access control policy. We set $V_0(\mathbf{x}) = \overline{U}^{\iota}_{\text{avg}}(\mathbf{x})$. The algorithm converges after 380 iterations to within $1e-4$. The feasible region and optimal coordinated connection access control policy are shown in Fig. 3. The cell can admit $\overline{x}_1 = 11$ class-1 applications, $\overline{x}_2 = 7$ class-2 applications, or a bit less than a linear combination of the two. In general, class-1 applications are preferred since they generate a higher average utility per unit rate at their minimum rate, i.e., $U_1(S_1')/S_1' > U_2(S_2')/S_2'$. When the state is on or above the upper boundary of the feasible region, the optimal CAC policy is to block all arrivals. When the state is far from the upper boundary, the optimal policy is to admit all arrivals. In a zone near the boundary as illustrated, the optimal policy is to admit only class-1 arrivals. As a result, the system will never enter a state with $x_2 > 4$.

In this scenario, it is worth noting that $\overline{U}^{\iota}_{\text{avg}}(\mathbf{x})$ is monotonically increasing with both x_1 and x_2 over the entire feasible region. If the feasible region was rectangular, then Theorem 1 would lead one to expect that the optimal policy would admit all arrivals while inside the feasible region. The region in which the optimal policy admits only class-1 arrivals comes from the nonrectangular shape of the feasible region.

We would like to compare the *coordinated connection access control* (CCAC) policy to the other two policies. The first alternate policy is a *traditional connection access control* (TCAC) policy that admits arrivals if and only if there are sufficient resources to guarantee QoS constraints

$$\forall \mathbf{x}, g_l = 0 \text{ if } \mathbf{x} + \mathbf{e}_l \notin \Lambda, \quad g_l = 1 \text{ otherwise}$$

i.e., that admits all users in all states inside the feasible region.

The second alternate policy is one that admits arrivals if and only if there are sufficient resources to guarantee QoS constraints and if the average utility in the state if admitted is higher than the average utility in the current state

$$\forall \mathbf{x},\ g_l = 0 \text{ if } \mathbf{x} + \mathbf{e}_l \notin \Lambda \text{ or if } \overline{U}^{\iota}_{\text{avg}}(\mathbf{x} + \mathbf{e}_l) < \overline{U}^{\iota}_{\text{avg}}(\mathbf{x})$$
$$g_l = 1 \text{ otherwise.}$$

We call this policy *greedy connection access control* (GCAC).

TABLE II
AVERAGE UTILITY FOR TWO SEMI-ELASTIC CLASSES

	TCAC	GCAC	CCAC
Utility	27.32	27.32	32.49
Blocking probability of type 1	64%	64%	32%
Blocking probability of type 2	56%	56%	81%
Total Utility of type 1	14.78	14.78	23.72
Total Utility of type 2	12.54	12.54	8.77

Clearly, the coordinated policy CCAC should perform at least as well as the greedy policy GCAC since it takes into account not only the average utility of the state if admitted, but also the expected utility earned in all future states. One may also expect that the greedy policy GCAC should outperform the traditional policy TCAC since it will block arrivals that decrease the current average utility $\overline{U}^{\iota}_{\text{avg}}(\mathbf{x})$.

To compare the CCAC, GCAC, and TCAC policies, we simulate 40 min of arrivals and departures; the resulting average utility under each policy is shown in Table II. In this scenario, arrivals always result in an increase in the current average utility $\overline{U}^{\iota}_{\text{avg}}(\mathbf{x})$, and thus the greedy policy GCAC is identical to the traditional policy TCAC. Under the TCAC policy and heavy load, the two application types experience similar blocking probabilities; the differences are due to the different bandwidth requirements of the applications and small differences in the discrete shape of the state space boundary. However, as noted above, the optimal coordinated policy CCAC blocks class-2 arrivals when the residual capacity is low. As a result, under CCAC the system spends a substantially greater portion of time near the highest utility state (11, 0) than it does under TCAC, and thus CCAC achieves a much higher average utility than TCAC and GCAC. The increase in utility is due to a reduction in the blocking probability for type-1 applications and a corresponding increase in blocking probability for type-2 applications. This may be viewed by some as unfair, reflecting the tension between fairness and efficiency discussed above.

We also examine how well each policy adheres to the QoS constraints. To do this, we focus on the outermost four slices, $m = 32$ to 35, which experience the lowest average rates. In Table III, we give the minimum average rate of all users in the given slice within each application class. We observe that each policy maintains average rates above the QoS requirement S_l' (300 kb/s for class 1, 625 kb/s for class 2) for all users.

Finally, we investigate the influence of the arrival rates of the two application classes on the connection access control policies. We consider four sets of arrival rates, of which two represent lighter loads and two heavier loads than used above. The optimal coordinated connection access control policy under each set of arrival rates is shown in Fig. 4. The feasible region remains the same, and class-1 applications remain preferred. When the load is very light (e.g., $\beta_1 = 1/20$, $\beta_2 = 1/40$) the optimal policy is to admit arrivals whenever the state upon admitting the arrival would remain in the state space, i.e., the TCAC policy. For light to heavy loads, as before there is a zone within the state space within which the optimal policy is to admit only class-1 arrivals. However, the location and size of this zone varies with the arrival rates. As the arrival rates increase, the

TABLE III
MINIMUM AVERAGE RATE (kb/s) FOR TWO SEMI-ELASTIC CLASSES

Slice number	TCAC		GCAC		CCAC	
	avg S_1^m	avg S_2^m	avg S_1^m	avg S_2^m	avg S_1^m	avg S_2^m
$m = 32$	504	756	504	756	529	753
$m = 33$	486	696	486	696	489	727
$m = 34$	445	661	445	661	476	729
$m = 35$	409	641	409	641	479	635

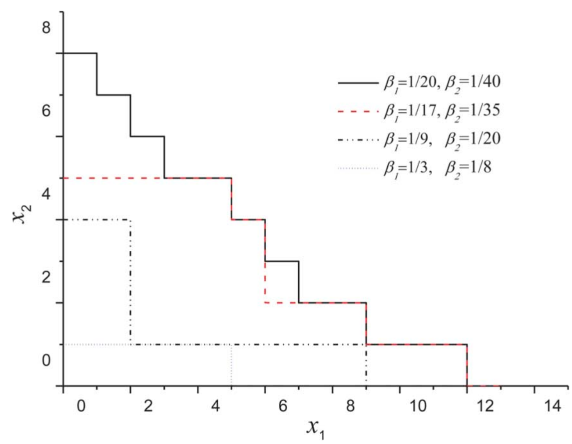

Fig. 4. Feasible region and CAC policy for two semi-elastic classes under various sets of arrival rates.

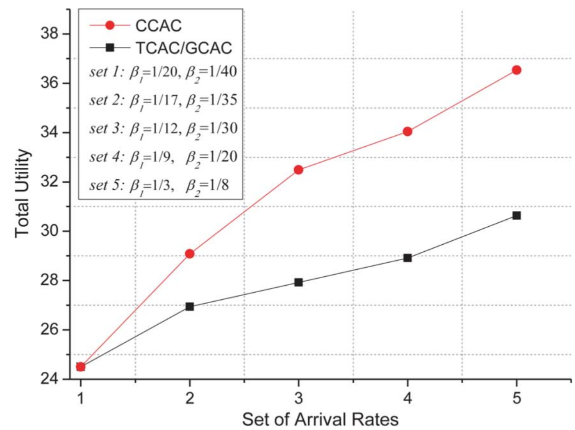

Fig. 5. Total utility under various sets of arrival rates.

zone expands, resulting in fewer class-2 arrivals being admitted. This occurs because the increased number of class-1 arrivals allows the system to be more selective about which arrivals to admit, as it becomes more likely that class-1 arrivals along will keep the system relatively busy.

In Fig. 5, the total utility under these four sets of arrival rates and the original set of arrival rates are illustrated for each policy. As before, the TCAC and GCAC policies are identical. The total utility is the sum of individual utilities that depend on the achieved rates. When arrival rates increase, although the average number of terms in the sum increases due to more active users, the utility per user decreases since more users compete for the same resources. A well-designed connection access policy ensures that the increase in the number of active users dominates the decrease in the utility per users. Here, all three policies are well designed in this sense since they block arrivals if it would place the state outside the feasible region. Therefore, the total utility increases with arrival rates under all three policies. However, the difference between the total utility earned by the CCAC policy and the other policies increases with the arrival rates, reflecting the increased benefit of the additional complexity of CCAC under heavier loads.

B. Elastic and Semi-Elastic Classes

We next focus on dynamic sharing of system resources between one class of elastic applications (e.g., FTP) and one class of semi-elastic applications (video). As above, we presume that 40% of system resources are allocated to these two classes, i.e., $N = 400$ subcarriers and $P = 42$ dbm downlink power, and that noise plus interference $\sigma^2 + I = 7.6232 * 10^{-28}$ mW. The time window over which utility is evaluated remains $W_2 =$ 133 slots for the video traffic, but is increased to $W_1 = 300$ slots for the FTP traffic since it is elastic. The class-1 application is elastic with utility function

$$U_1(S_{k,t^\iota}) = 1.2 * (S_{k,t^\iota} + 1)^{\frac{1}{3.15}} - 1.2$$

where S_{k,t^ι} is expressed in units of 100 kb/s. The class-2 application is semi-elastic with the same utility function as the class-1 application in Section V-A. The arrival processes are Poisson processes with intensities $\beta_1 = 1/5/s$ and $\beta_2 = 1/15/s$, respectively, and the connection durations are exponentially distributed with means $1/\eta_1 = 600$ s and $1/\eta_2 = 180$ s, respectively.

Because there is no minimum rate requirement for elastic applications, there is no hard limit on the total number of active users running elastic applications. However, we observe that deployed systems often limit the number of active data users per cell, and thus we set an upper limit $\overline{x}_1 = 200$. The limit is assumed to be independent of the number of active video users, and thus the feasible region is rectangular. The limit on the number of active video users remains $\overline{x}_2 = 11$ as in Section V-A. The value iteration algorithm converges after 1890 iterations to within $1e - 3$. The feasible region is shown in Fig. 6.

The average utility $\overline{U}^\iota_{\text{avg}}(\mathbf{x})$ is monotonically increasing with x_1 for all fixed x_2. Thus by Theorem 1, the optimal policy always admits class-1 arrivals.

However, the average utility $\overline{U}^\iota_{\text{avg}}(\mathbf{x})$ is not monotonic with x_2 for all fixed x_1, and thus the optimal policy does not always admit class-2 arrivals. In this scenario, the feasible region can be partitioned as illustrated in Fig. 7; in the region below the curve, $\overline{U}^\iota_{\text{avg}}(\mathbf{x})$ is monotonically increasing with x_2 for fixed

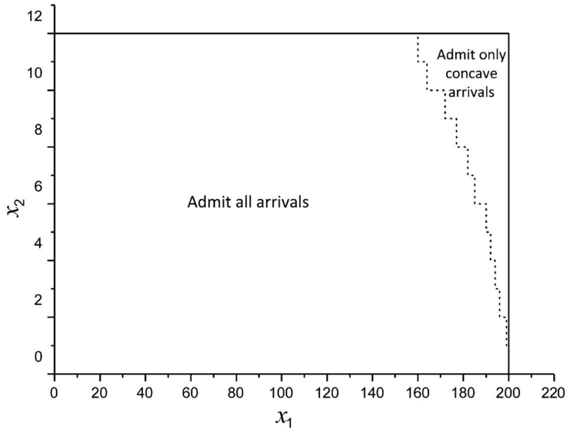

Fig. 6. Feasible region and CAC for elastic and semi-elastic classes.

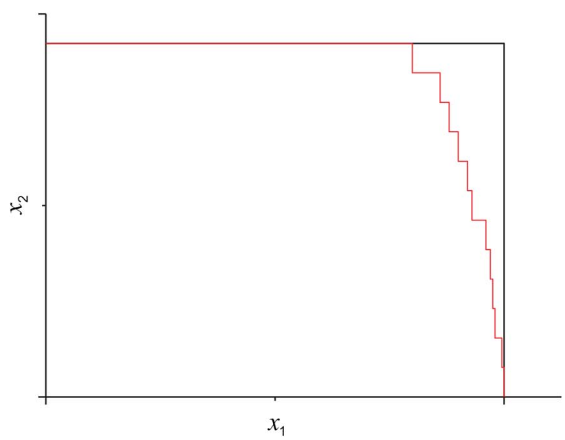

Fig. 7. Regions of monotonicity.

TABLE IV
AVERAGE UTILITY FOR ELASTIC AND SEMI-ELASTIC CLASSES

	TCAC	GCAC	CCAC
Utility	89.73	94.42	98.42

x_1, and in the region above the curve, $\overline{U}^{\iota}_{\text{avg}}(\mathbf{x})$ is monotonically decreasing with x_2 for fixed x_1. The optimal coordinated connection access control policy is shown in Fig. 6. It admits all class-1 arrivals while inside the feasible region (as guaranteed by Theorem 1), but only admits class-2 arrivals when the state is below the threshold illustrated. This optimal admission threshold is different than the partition based on the average utility $\overline{U}^{\iota}_{\text{avg}}(\mathbf{x})$. A more detailed investigation of the form of the optimal policy is presented in Appendix III.

To compare the CCAC, GCAC, and TCAC policies, we simulate 30 min of arrivals and departures; the resulting average utility under each policy is shown in Table IV. Because class-2 arrivals result in a decrease in the current average utility $\overline{U}^{\iota}_{\text{avg}}(\mathbf{x})$ when the state is above the curve in Fig. 7, the greedy policy GCAC outperforms the traditional policy TCAC. However, the coordinated policy CCAC outperforms GCAC by placing the threshold at a slight different location than the curve.

Finally, we examine how well each policy adheres to the QoS constraints. As above, we focus on the outermost four slices, $m = 32$ to 35, which experience the lowest average rates. Recall that the elastic class-1 applications do not have a QoS requirement. In Table V, we give the minimum average rate of all class-2 users in the given slice. Unlike the scenario with two semi-elastic classes, in this scenario we observe that users in the furthest slice from the base station experience an average rate below their QoS requirement (300 kb/s) under each of the three policies. This causes outage for some users far from the base station. Such outage can be addressed by adding an outage constraint that can be efficiently enforced by adding an outage price to the rate price; see [37] for details.

TABLE V
MINIMUM AVERAGE RATE (IN kb/s) FOR ELASTIC AND SEMI-ELASTIC CLASSES

Slice number	Avg S_2^m of TCAC	Avg S_2^m of GCAC	Avg S_2^m of CCAC
$m = 32$	346	331	353
$m = 33$	325	326	313
$m = 34$	313	306	305
$m = 35$	296	286	287

C. Elastic and Inelastic Classes

Our last scenario focuses on dynamic sharing of system resources between one class of elastic applications (e.g., FTP) and one class of inelastic applications (voice). Since these are two popular classes, we presume that 80% of system resources are allocated to these two classes, i.e., $N = 800$ subcarriers and $P = 45$ dbm downlink power, and that noise plus interference $\sigma^2 + I = 7.8288 * 10^{-28}$ mW. The time window over which utility is evaluated remains $W_1 = 300$ slots for the FTP traffic and is set to a short $W_2 = 10$ slots for the voice traffic since it is inelastic. The class-1 application is elastic with the same utility function given in Section V-B. The class-2 application is inelastic with utility function

$$U_2(S_{k,t^{\iota}}) = \begin{cases} 1.5, & \text{if } S_{k,t^{\iota}} > 32 \text{ kb/s} \\ 0, & \text{else} \end{cases}$$

where $S_{k,t^{\iota}}$ is expressed in units of 100 kb/s. The arrival processes are Poisson processes with intensities $\beta_1 = 1/4/s$ and $\beta_2 = 1/3/s$ respectively, and the connection durations are Exponentially distributed with means $1/\eta_1 = 100$ s and $1/\eta_2 = 300$ s, respectively.

As in the previous scenario, because there is no minimum rate requirement for elastic applications, we set an upper limit $\overline{x}_1 = 200$, independent of the number of active voice users, and thus the feasible region is rectangular. The limit on the number of active voice users remains $\overline{x}_2 = 140$. The value iteration algorithm converges after 5160 iterations to within $1e-3$. The feasible region is shown in Fig. 8.

The average utility $\overline{U}^{\iota}_{\text{avg}}(\mathbf{x})$ is monotonically increasing with both x_1 and x_2, and thus the optimal policy admits all arrivals while in the feasible region. The coordinated policy CCAC is thus identical to the TCAC and GCAC policies. The resulting average utility (from a 30-min simulation) is shown in Table VI.

Finally, we examine how well each policy adheres to the QoS constraints. As above, we focus on the outermost four slices,

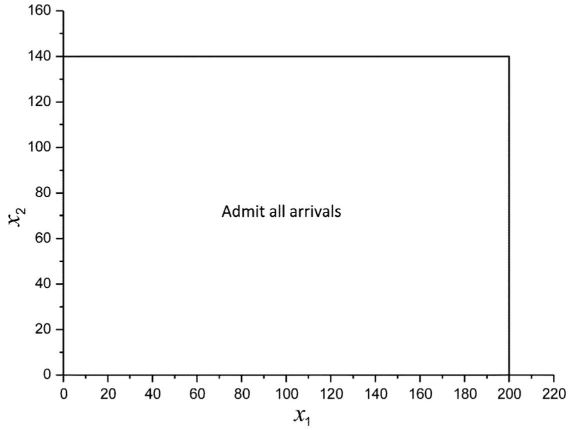

Fig. 8. Feasible region and CAC for semi-elastic and inelastic classes.

TABLE VI
AVERAGE UTILITY FOR ELASTIC AND INELASTIC CLASSES

	TCAC	GCAC	CCAC
Utility	237.28	237.28	237.28

TABLE VII
MINIMUM AVERAGE RATE (IN kb/s) FOR ELASTIC AND INELASTIC CLASSES

Slice number	Avg S_2^m of TCAC	Avg S_2^m of GCAC	Avg S_2^m of CCAC
$m = 32$	32.9	32.9	32.9
$m = 33$	32.7	32.7	32.7
$m = 34$	32.2	32.2	32.2
$m = 35$	32.3	32.3	32.3

$m = 32$ to 35, which experience the lowest average rates. Recall that the elastic class-1 applications do not have a QoS requirement. In Table VII, we give the minimum average rate of all class-2 users in the given slice. We observe the policy maintains average rates above the QoS requirement of 32 kb/s for voice users.

VI. CONCLUSION

We have investigated how connection access control and resource allocation can be coordinated so that both focus on utility. Whereas previous resource allocation proposals have principally focused on maximizing instantaneous or short-term throughput or utility, and connection access control algorithms often simply admit users if there is sufficient capacity, here we have proposed that CAC and RA should both focus on maximizing long-term average utility.

We have found that joint CAC and RA can be decomposed into separate CAC and RA modules that treat each other as black boxes, and that the interface simply consists of an exchange of information about feasible states, current state, and average utilities in each state. This allows the RA module to maximize the long-term average utility for each system state, and the CAC module to maximize the long-term average utility over all system states. This decomposition provides flexibility in the design of the CAC and RA policies.

We also gave examples of RA and CAC algorithms that satisfy this decomposition and interface. The RA module determines a set of rate prices for each application based on a user's combined slow fading, shadowing, and path loss and uses these rate prices to allocate power and subcarriers. The CAC module uses stochastic dynamic programming to determine the optimal admission policy.

Simulation showed that the optimal CAC policy may block applications with relatively low average utility per unit rate even when capacity is available. This preferential treatment by application can result in much higher average utility while continuing to satisfy QoS constraints.

APPENDIX I

This appendix provides the derivation of (5). The first-order conditions for (4) result in an optimal power allocation

$$p_{k,n,t^\iota} = \left(\frac{B\overline{\lambda}_{k,t^\iota}}{\mu_{t^\iota} \ln 2} - \frac{\sigma^2 + I}{|H_{k,n,t^\iota}|^2} \right)^+ \tag{10}$$

and that subcarrier n should be allocated to the user

$$\arg\max_k \Phi_{k,n,t^\iota} \tag{11}$$

where

$$\Phi_{k,n,t^\iota} = \overline{\lambda}_{k,t^\iota} B \left[\log_2 \left(c \frac{B\overline{\lambda}_{k,t^\iota}}{\mu_{t^\iota} \ln 2} \frac{|H_{k,n,t^\iota}|^2}{\sigma^2 + I} \right) \right]^+ - \mu_{t^\iota} \left(\frac{B\overline{\lambda}_{k,t^\iota}}{\mu_{t^\iota} \ln 2} - \frac{\sigma^2 + I}{|H_{k,n,t^\iota}|^2} \right)^+$$

where $\overline{\lambda}_{k,t^\iota} = \sum_{\tau^\iota = t^\iota}^{t^\iota + W - 1} \lambda_{k,\tau^\iota} / W_k$. The bound on the achieved rate is

$$d_{k,t^\iota} = \max \left\{ S'_{b_k}, \arg\max_{d_{k,t^\iota}} \left[U_{b_k}(d_{k,t^\iota})/T^\iota - \overline{\lambda}_{k,t^\iota} d_{k,t^\iota} \right] \right\}. \tag{12}$$

The fluctuations in fast fading will average out during time window W_{b_k}, whereas fluctuations in slow fading, shadowing, and path loss will not. Thus, we propose basing the average rate price $\overline{\lambda}_{k,t^\iota}$ on a user's combined path loss and shadowing, i.e., $\overline{\lambda}_{k,t^\iota}$ should be a function of ψ_k, denoted $\lambda_k(\psi_k)$. Furthermore, we propose using this average price to determine the bound d_{k,t^ι} by replacing λ_{k,t^ι} in (12) $\overline{\lambda}_{k,t^\iota}$ so that they vary on the same timescale.

User k's achieved rate S_k depends not only on its ψ_k, but also upon other users' $\boldsymbol{\psi}_{-k}$ and on all users' fast fading. Since fast fading will largely average out within time window W_{b_k}, the expectation over $\boldsymbol{\alpha^2}$ can be brought inside the utility function without loss of accuracy. Other users' shadowing and path losses affect the resource allocation through determination of the power price μ_{k,t^ι}, but consideration of them in determination of the resource allocation policy $\{\lambda_k(\psi_k),\ \forall k\}$ is too complex; thus, we also bring $\boldsymbol{\psi}_{-k}$ inside the utility function. Then

$$\sum_{k=1}^{K} E_{\boldsymbol{\alpha^2},\boldsymbol{\psi}} U_{b_k}(S_k) \approx \sum_{k=1}^{K} E_{\psi_k} U_k(E_{\boldsymbol{\alpha^2},\boldsymbol{\psi}_{-k}} S_k).$$

Similarly, the total power can be written as a statistical average over an infinite time horizon

$$\frac{1}{T^\iota}\sum_{t^\iota=1}^{T^\iota}\sum_{k=1}^{K}\sum_{n=1}^{N} p_{k,n,t^\iota} \to \sum_{k=1}^{K} E_{\psi_k}\sum_{n=1}^{N} E_{\boldsymbol{\alpha^2},\boldsymbol{\psi}_{-k}}(p_{k,n}).$$

The QoS requirement can be written as

$$S_{k,t^\iota} > S'_{b_k}\ \forall k, t^\iota \to E_{\boldsymbol{\alpha^2},\boldsymbol{\psi}_{-k}} S_k(\psi_k) > S'_{b_k},\ \forall k, \psi_k.$$

We partition ψ_k into M slices. This leads to the quantized optimization problem (5).

Appendix II

This appendix gives the proof for Theorem 1.

Lemma: If $V_i(\mathbf{x})$ is monotonically increasing with x_1 for fixed x_2, then $V_{i+1}(\mathbf{x})$ is monotonically increasing with x_1 for fixed x_2.

Proof: Using (9), we can express $V_{i+1}(\mathbf{x})$ as

$$\begin{aligned} V_{i+1}(\mathbf{x}) = {} & V_i(\mathbf{x}) + \overline{U}^\iota_{\text{avg}}(\mathbf{x}) \\ & + \underbrace{\nu x_1\eta_1\left[V_i(\mathbf{x}-\mathbf{e}_1) - V_i(\mathbf{x})\right]}_{a} \\ & + \underbrace{\nu x_2\eta_2\left[V_i(\mathbf{x}-\mathbf{e}_2) - V_i(\mathbf{x})\right]}_{b} \\ & + \mathbf{1}\left(V_i(\mathbf{x}+\mathbf{e}_1) > V_i(\mathbf{x})\right) \\ & \underbrace{\cdot\, \nu\beta_1\left[V_i(\mathbf{x}+\mathbf{e}_1) - V_i(\mathbf{x})\right]}_{c} \\ & + \mathbf{1}\left(V_i(\mathbf{x}+\mathbf{e}_2) > V_i(\mathbf{x})\right) \\ & \underbrace{\cdot\, \nu\beta_2\left[V_i(\mathbf{x}+\mathbf{e}_2) - V_i(\mathbf{x})\right]}_{d} \end{aligned}$$

where the terms including $V_i(\mathbf{x}-\mathbf{e}_l)$ are included only if $x_l > 0$ and the terms including $V_i(\mathbf{x}+\mathbf{e}_l)$ are included only if $\mathbf{x}+\mathbf{e}_l \in \Lambda^\iota$. Similarly, denote the corresponding terms in $V_{i+1}(\mathbf{x}+\mathbf{e}_1)$ by a', b', c', d'.

We wish to show that if $V_i(\mathbf{x}+\mathbf{e}_1) - V_i(\mathbf{x}) > 0$, then $V_{i+1}(\mathbf{x}+\mathbf{e}_1) - V_{i+1}(\mathbf{x}) > 0$. By the hypothesis, the terms a and a' are negative. Also, the terms c, c', d, and d' are positive if they exist. Consider $V_{i+1}(\mathbf{x}+\mathbf{e}_1) - V_{i+1}(\mathbf{x})$ in two cases: $x_2 > 0$ and $x_2 = 0$.

Case 1: $x_2 > 0$: The number of terms in the difference depends on whether terms a, b, c, and/or d are included in $V_{i+1}(\mathbf{x})$, and similarly whether terms a', b', c', and/or d' are included in $V_{i+1}(\mathbf{x}+\mathbf{e}_1)$. Since $x_2 > 0$, $V_{i+1}(\mathbf{x})$ includes term b, and thus $V_{i+1}(\mathbf{x}+\mathbf{e}_1)$ includes term b'. We will examine the situation in which $V_{i+1}(\mathbf{x})$ contain only terms b, c and d, and $V_{i+1}(\mathbf{x}+\mathbf{e}_1)$ contains only terms a' and term b', i.e., $V_{i+1}(\mathbf{x})$ only includes positive terms c and d, and $V_{i+1}(\mathbf{x}+\mathbf{e}_1)$ only includes negative term a. In this situation, the difference is

$$\begin{aligned} & V_{i+1}(\mathbf{x}+\mathbf{e}_1) - V_{i+1}(\mathbf{x}) \\ & = \left[1-\nu\left((x_1+1)\eta_1 + x_2\eta_2 + \beta_1\right)\right]\left[V_i(\mathbf{x}+\mathbf{e}_1) - V_i(\mathbf{x})\right] \\ & \quad + \left[\overline{U}^\iota_{\text{avg}}(\mathbf{x}+\mathbf{e}_1) - \overline{U}^\iota_{\text{avg}}(\mathbf{x})\right] \\ & \quad + \nu x_2\eta_2\left[V_i(\mathbf{x}+\mathbf{e}_1-\mathbf{e}_2) - V_i(\mathbf{x}-\mathbf{e}_2)\right] \\ & \quad - \nu\beta_2\left[V_i(\mathbf{x}+\mathbf{e}_2) - V_i(\mathbf{x})\right]. \end{aligned}$$

Since $V_{i+1}(\mathbf{x})$ contains term d and $V_{i+1}(\mathbf{x}+\mathbf{e}_1)$ does not contain term d', it means $V_i(\mathbf{x}+\mathbf{e}_1) > V_i(\mathbf{x}+\mathbf{e}_1+\mathbf{e}_2) > V_i(\mathbf{x}+\mathbf{e}_2)$. Thus

$$\begin{aligned} & V_{i+1}(\mathbf{x}+\mathbf{e}_1) - V_{i+1}(\mathbf{x}) \\ & \geq \left[1-\nu\left((x_1+1)\eta_1 + x_2\eta_2 + \beta_1\right)\right]\left[V_i(\mathbf{x}+\mathbf{e}_1) - V_i(\mathbf{x})\right] \\ & \quad + \left[\overline{U}^\iota_{\text{avg}}(\mathbf{x}+\mathbf{e}_1) - \overline{U}^\iota_{\text{avg}}(\mathbf{x})\right] \\ & \quad + \nu x_2\eta_2\left[V_i(\mathbf{x}+\mathbf{e}_1-\mathbf{e}_2) - V_i(\mathbf{x}-\mathbf{e}_2)\right] \\ & \quad - \nu\beta_2\left[V_i(\mathbf{x}+\mathbf{e}_1) - V_i(\mathbf{x})\right] \\ & = \left[1-\nu\left((x_1+1)\eta_1 + x_2\eta_2 + \beta_1 + \beta_2\right)\right]\left[V_i(\mathbf{x}+\mathbf{e}_1) - V_i(\mathbf{x})\right] \\ & \quad + \left[\overline{U}^\iota_{\text{avg}}(\mathbf{x}+\mathbf{e}_1) - \overline{U}^\iota_{\text{avg}}(\mathbf{x})\right] \\ & \quad + \nu x_2\eta_2\left[V_i(\mathbf{x}+\mathbf{e}_1-\mathbf{e}_2) - V_i(\mathbf{x}-\mathbf{e}_2)\right]. \end{aligned}$$

Each term in brackets is greater than zero, and thus $V_{i+1}(\mathbf{x}+\mathbf{e}_1) - V_{i+1}(\mathbf{x}) > 0$. In situations where $V_{i+1}(\mathbf{x})$ and/or $V_{i+1}(\mathbf{x}+\mathbf{e}_1)$ contain different terms, it can be similarly shown that $V_{i+1}(\mathbf{x}+\mathbf{e}_1) - V_{i+1}(\mathbf{x}) > 0$.

Case 2: $x_2 = 0$: We will examine the situation in which $V_{i+1}(\mathbf{x})$ contains only terms c and d, and $V_{i+1}(\mathbf{x}+\mathbf{e}_1)$ contains only term a'. In this situation, the difference is

$$\begin{aligned} & V_{i+1}(\mathbf{x}+\mathbf{e}_1) - V_{i+1}(\mathbf{x}) \\ & = \left[1-\nu\left((x_1+1)\eta_1 + \beta_1\right)\right]\left[V_i(\mathbf{x}+\mathbf{e}_1) - V_i(\mathbf{x})\right] \\ & \quad + \left[\overline{U}^\iota_{\text{avg}}(\mathbf{x}+\mathbf{e}_1) - \overline{U}^\iota_{\text{avg}}(\mathbf{x})\right] - \nu\beta_2\left[V_i(\mathbf{x}+\mathbf{e}_2) - V_i(\mathbf{x})\right] \\ & \geq \left[1-\nu\left((x_1+1)\eta_1 + \beta_1\right)\right]\left[V_i(\mathbf{x}+\mathbf{e}_1) - V_i(\mathbf{x})\right] \\ & \quad + \left[\overline{U}^\iota_{\text{avg}}(\mathbf{x}+\mathbf{e}_1) - \overline{U}^\iota_{\text{avg}}(\mathbf{x})\right] - \nu\beta_2\left[V_i(\mathbf{x}+\mathbf{e}_1) - V_i(\mathbf{x})\right] \\ & = \left[1-\nu\left((x_1+1)\eta_1 + \beta_1 + \beta_2\right)\right]\left[V_i(\mathbf{x}+\mathbf{e}_1) - V_i(\mathbf{x})\right] \\ & \quad + \left[\overline{U}^\iota_{\text{avg}}(\mathbf{x}+\mathbf{e}_1) - \overline{U}^\iota_{\text{avg}}(\mathbf{x})\right]. \end{aligned}$$

Each term in brackets is greater than zero, and thus $V_{i+1}(\mathbf{x}+\mathbf{e}_1) - V_{i+1}(\mathbf{x}) > 0$. In situations where $V_{i+1}(\mathbf{x})$ and/or $V_{i+1}(\mathbf{x}+\mathbf{e}_1)$ contain different terms, it can be similarly shown that $V_{i+1}(\mathbf{x}+\mathbf{e}_1) - V_{i+1}(\mathbf{x}) > 0$. ■

We can prove Theorem 1 by mathematical induction. For the base case, we set $V_0(\mathbf{x}) = \overline{U}^\iota_{\text{avg}}(\mathbf{x})$. Since, $\overline{U}^\iota_{\text{avg}}(\mathbf{x})$ is monotonically increasing with x_1 for fixed x_2, so is $V_0(\mathbf{x})$. The induction step has been proven by Lemma 1. Thus, $V_\infty(\mathbf{x})$ is monotonically increasing with x_1 for fixed x_2. It follows that the corresponding CAC policy always admits class-1 arrivals while inside the feasible region.

Appendix III

This appendix explores the form of the optimal policy for the scenario in Section V for elastic and semi-elastic classes. To gain some insight, consider a rectangular feasible region $\{(x_1, x_2)$ such that $0 \leq x_1 \leq \bar{x}_1, 0 \leq x_2 \leq \bar{x}_2\}$ on which average utility $\overline{U}^\iota_{\text{avg}}(\mathbf{x})$ is: 1) monotonically increasing with x_1; 2) monotonically increasing with x_2 on $0 \leq x_2 \leq a$; and 3) monotonically decreasing with x_2 on $a < x_2 \leq \bar{x}_2$. This example satisfies the hypothesis for Theorem 1, and thus the CAC policy that satisfies (9) always admits class-1 arrivals while inside the feasible region.

It is straightforward to prove that $V_i(\mathbf{x})$ is monotonically increasing with x_2 on $0 \leq x_2 \leq a$, and thus the iteration i CAC policy also admits class-2 arrivals on $0 \leq x_2 < a$ while inside the feasible region.

It remains to determine the iteration i policy on $a \leq x_2 < \bar{x}_2$. On $0 < x_1 < \bar{x}_1$, $a < x_2 < \bar{x}_2$

$$\begin{aligned} &V_{i+1}(\mathbf{x}+\mathbf{e}_2) - V_{i+1}(\mathbf{x}) \\ &= \left[1-\nu\left(x_1\eta_1+(x_2+1)\eta_2+\beta_1\right)\right]\left[V_i(\mathbf{x}+\mathbf{e}_2)-V_i(\mathbf{x})\right] \\ &\quad + \left[\overline{U}^{\iota}_{\mathrm{avg}}(\mathbf{x}+\mathbf{e}_2) - \overline{U}^{\iota}_{\mathrm{avg}}(\mathbf{x})\right] \\ &\quad + \nu x_1\eta_1\left[V_i(\mathbf{x}+\mathbf{e}_2-\mathbf{e}_1)-V_i(\mathbf{x}-\mathbf{e}_1)\right] \\ &\quad + \nu x_2\eta_2\left[V_i(\mathbf{x})-V_i(\mathbf{x}-\mathbf{e}_2)\right] \\ &\quad + \nu\beta_1\left[V_i(\mathbf{x}+\mathbf{e}_2+\mathbf{e}_1)-V_i(\mathbf{x}+\mathbf{e}_1)\right]. \end{aligned} \tag{13}$$

Each term in brackets is negative, and thus the iteration i policy blocks class-2 arrivals in this subset of the state space.

On the boundary between the two regions of monotonicity, it is a more complex situation. On $0 < x_1 < \bar{x}_1$ and $x_2 = a$, the value iteration can still be written as (13). However, the term $\nu x_2\eta_2[V_i(\mathbf{x}) - V_i(\mathbf{x}-\mathbf{e}_2)]$ is now positive. The iteration i policy may thus either admit or block class-2 arrivals on the boundary between the two regions of monotonicity. As the iteration number increases, the boundary between the two regions of monotonicity may also move. Thus, while the optimal policy is likely to be a threshold policy, the threshold is likely to change from the initial boundary.

In the scenario with one elastic and one semi-elastic class, we see a similar phenomenon. The average utility $\overline{U}^{\iota}_{\mathrm{avg}}(\mathbf{x})$ is: 1) monotonically increasing with x_1; 2) monotonically increasing with x_2 below the curve; and 3) monotonically decreasing with x_2 above the curve. The optimal policy is a threshold policy, but the threshold differs slightly from the curve.

REFERENCES

[1] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication.* . Cambridge, U.K.: Cambridge Univ. Press, 2005.

[2] J. Jang and K. B. Lee, "Transmit power adaptation for multiuser OFDM systems," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 2, pp. 171–178, Feb. 2003.

[3] Z. Shen, J. Andrews, and B. Evans, "Adaptive resource allocation in multiuser OFDM systems with proportional rate constraints," *IEEE Trans. Wireless Commun.*, vol. 4, no. 6, pp. 2726–2737, Nov. 2005.

[4] M. Tao, Y.-C. Liang, and F. Zhang, "Resource allocation for delay differentiated traffic in multiuser OFDM systems," *IEEE Trans. Wireless Commun.*, vol. 7, no. 6, pp. 2190–2201, Jun. 2008.

[5] I. Wong and B. Evans, "Optimal resource allocation in the OFDMA downlink with imperfect channel knowledge," *IEEE Trans. Commun.*, vol. 57, no. 1, pp. 232–241, Jan. 2009.

[6] P. Liu, P. Zhang, S. Jordan, and M. Honig, "Single-cell forward link power allocation using pricing in wireless networks," *IEEE Trans. Wireless Commun.*, vol. 3, no. 2, pp. 533–543, Mar. 2004.

[7] G. Song and Y. Li, "Utility-based resource allocation and scheduling in OFDM-based wireless broadband networks," *IEEE Commun. Mag.*, vol. 43, no. 12, pp. 127–134, Dec. 2005.

[8] C. Yang, W. Wang, and X. Zhang, "Multi-service transmission in multiuser cooperative networks," in *Proc. WCNC*, Apr. 2009, pp. 1–5.

[9] C. Yang and S. Jordan, "Resource allocation for semi-elastic applications in wireless networks," in *Proc. IEEE Global Telecommun. s Conf.*, Dec. 2011, pp. 1–6.

[10] Z. Liu and M. El Zarki, "SIR-based call admission control for DS-CDMA cellular systems," *IEEE J. Sel. Areas Commun.*, vol. 12, no. 4, pp. 638–644, May 1994.

[11] D. Kim, "Efficient interactive call admission control in power-controlled mobile systems," *IEEE Trans. Veh. Technol.*, vol. 49, no. 3, pp. 1017–1028, May 2000.

[12] S. Anand and A. Chockalingam, "Performance analysis of voice/data cellular CDMA with SIR-based admission control," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 10, pp. 1674–1684, Sep. 2003.

[13] S. Bashar and Z. Ding, "Admission control and resource allocation in a heterogeneous OFDMA wireless network," *IEEE Trans. Wireless Commun.*, vol. 8, no. 8, pp. 4200–4210, Aug. 2009.

[14] X. Gong, S. Vorobyov, and C. Tellambura, "Joint bandwidth and power allocation with admission control in wireless multi-user networks with and without relaying," *IEEE Trans. Signal Process,*, vol. 59, no. 4, pp. 1801–1813, Apr. 2011.

[15] A. Goldsmith and S.-G. Chua, "Adaptive coded modulation for fading channels," *IEEE Trans. Commun.*, vol. 46, no. 5, pp. 595–602, May 1998.

[16] I. S. Gopal and T. E. Stem, "Optimal call blocking policies in an integrated services environment," in *Proc. Conf Inf. Sci. Sysf.*, 1983, pp. 383–388.

[17] K. Miyake, "Optimal trunk reservation control for multislot connection circuits," *Trans. IEICE*, vol. J71-B, no. 12, pp. 1419–1424, Nov. 1988.

[18] K. Ross and D. H. K. Tsang, "Optimal circuit access policies in an ISDN environment: A markov decision approach," *IEEE Trans. Commun.*, vol. 37, no. 9, pp. 934–939, Sep. 1989.

[19] T. Oda and Y. Watanabe, "Optimal trunk reservation for a group with multislot traffic streams," *IEEE Trans. Commun.*, vol. 38, no. 7, pp. 1078–1084, Jul. 1990.

[20] G. R. Ash, J. Chen, A. E. Frey, and B. D. Huang, "Real-time network routing in a dynamic class-of-service network," in *Proc. ITC 13*, Jun. 1991.

[21] A. Gersht and K. Lee, "Virtual-circuit load control in fast packet-switched broadband networks," in *Proc. IEEE Global Telecommun. s Conf.*, 1988, vol. 1, pp. 214–220.

[22] Z. Dziong and L. Mason, "Fair-efficient call admission control policies for broadband networks-a game theoretic framework," *IEEE/ACM Trans. Netw.*, vol. 4, no. 1, pp. 123–136, Feb. 1996.

[23] C. Y. Wong, R. Cheng, K. Lataief, and R. Murch, "Multiuser OFDM with adaptive subcarrier, bit, power allocation," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 10, pp. 1747–1758, Oct. 1999.

[24] L. Tassiulas and S. Sarkar, "Maxmin fair scheduling in wireless networks," in *Proc. 21st Annu. IEEE INFOCOM*, 2002, vol. 2, pp. 763–772.

[25] A. Sridharan and B. Krishnamachari, "Max-min fair collision-free scheduling for wireless sensor networks," in *Proc. IEEE Int. Conf. Perform., Comput., Commun.*, 2004, pp. 585–590.

[26] A. Jalali, R. Padovani, and R. Pankaj, "Data throughput of CDMA-HDR a high efficiency-high data rate personal communication wireless system," in *Proc. IEEE Veh. Technol. Conf.*, 2000, vol. 3, pp. 1854–1858.

[27] P. Viswanath, D. Tse, and R. Laroia, "Opportunistic beamforming using dumb antennas," *IEEE Trans. Inf. Theory,* , vol. 48, no. 6, pp. 1277–1294, Jun. 2002.

[28] *E-UTRA Radio Frequency (RF) System Scenarios*, 3GPP TR 36.942 v10.2.0 Std., May 2011.

[29] C. Yang and S. Jordan, "Power and rate allocation for video conferencing in cellular networks," in *Proc. 49th Annu. Allerton Conf. Commun., Control, Comput.*, Sep. 2011, pp. 127–134.

[30] H. C. Tijms, *Stochastic Modelling and Analysis: A Computational Approach.*. New York, NY, USA: Wiley, 1986.

[31] G. Foschini, B. Gopinath, and J. Hayes, "Optimum allocation of servers to two types of competing customers," *IEEE Trans. Commun.*, vol. COM-29, no. 7, pp. 1051–1055, Jul. 1981.

[32] K. Ross and D. Tsang, *Stochastic Knapsack Problem*, vol. 37, no. 7, pp. 740–747, Jul. 1989.

[33] E. Altman, T. Jimenez, and G. Koole, "On optimal call admission control in resource-sharing system," *IEEE Trans. Commun.*, vol. 49, no. 9, pp. 1659–1668, Sep. 2001.

[34] S. Jordan and P. Varaiya, "Control of multiple service, multiple resource communication networks," *IEEE Trans. Commun.*, vol. 42, no. 11, pp. 2979–2988, Nov. 1994.

[35] *Evolved Universal Terrestrial Radio Access*, 3GPP TR 36.300 v10.2.0 Std., Jan. 2011.

[36] *Selection Procedures for the Choice of Radio Transmission Technologies of the UMTS*, 3GPP TR 30.03 Std., Apr. 1998.

[37] C. Yang and S. Jordan, "Resource allocation for semi-elastic applications with outage constraints in cellular networks," in *Proc. Asilomar Conf. Signals, Syst. Comput.*, Nov. 2013, pp. 1176–1182.

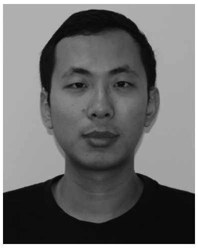

Chao Yang received the B.S. and M.S. degrees in electrical engineering from Beijing University of Posts and Telecommunications, Beijing, China, in 2006 and 2009, respectively, and the Ph.D. degree in networked systems from the University of California, Irvine, CA, USA, in 2014.

His research interests include resource allocation, admission control, and network optimization for both computer networks and cellular networks.

Scott Jordan (S'86–M'90) received the A.B. degree in applied mathematics and the B.S., M.S., and Ph.D. degrees in electrical engineering and computer science from the University of California, Berkeley, in 1985, 1987, and 1990, respectively.

From 1990 until 1999, he served as a faculty member with Northwestern University, Evanston, IL, USA, Since 1999, he has served as a faculty member with the University of California, Irvine, CA, USA. During 2006, he served as an IEEE Congressional Fellow, working in the United States Senate on Internet and telecommunications policy issues. His research interests currently include net neutrality, pricing and differentiated services in the Internet, and resource allocation in wireless multimedia networks.