On-line Pricing of Secondary Spectrum Access with Unknown Demand Function and Call Length Distribution

Huseyin Mutlu, Murat Alanyali, and David Starobinski
Department of Electrical and Computer Engineering
Boston University, Boston, MA 02215
E-mail: {hmutlu,alanyali,staro}@bu.edu

***Abstract*—We consider a wireless provider who caters to two classes of customers, namely primary and secondary users. Primary users have long term contracts while secondary users are admitted and priced according to current availability of excess spectrum. Secondary users accept an advertised price with a certain probability defined by an underlying demand function. We analyze the problem of maximizing profit gained by admission of secondary users. Previous studies in the field usually assume that the demand function is known and that the call length distribution is also known and exponentially distributed. In this paper, we analyze more realistic settings where both of these quantities are unknown. Our main contribution is to derive near-optimal pricing strategies under such settings. We focus on occupancy-based pricing policies, which depend only on the total number of ongoing calls in the system. We first show that such policies are insensitive to call length distribution except through the mean. Next, we introduce a new on-line, occupancy-based pricing algorithm, called Measurement-based Threshold Pricing (MTP) that operates by measuring the reaction of secondary users to a specific price and does not require the demand function to be known. MTP optimizes a profit function that depends on price only. We prove that while the profit function can be multimodal, MTP converges to one of the local optima as fast as if the function were unimodal. Lastly, we provide numerical studies demonstrating the near-optimal performance of occupancy-based policies for diverse sets of call length distributions and demand functions and the quick convergence of MTP to near-optimal on-line profit.**

I. Introduction

As a result of continuing efforts to deregulate wireless spectrum management, policy agencies are granting providers with the right to lease their spectrum [1]. This policy reform promises more efficient use of excess spectrum, which otherwise may be wasted. Implications of this reform can be seen in the novel services provided by spectrum brokerage companies, which match potential lessees and spectrum providers (licence holders). One such service is an on-line spectrum trading and leasing platform, called SpecEx.com, which was launched by Spectrum Bridge Inc. in 2008 [2].

Success of the aforementioned spectrum reforms hinges on the design of efficient pricing strategies, since a spectrum provider strives to maximize its profit from leasing its excess spectrum. In this paper, we aim at developing a realistic pricing framework to achieve this goal. We consider a wireless spectrum provider who caters to two classes of customers namely primary users (PU) and secondary users (SU). PUs have long term contracts and are not subjected to on-line pricing. On the other hand, SUs are admitted and priced according to current availability of excess spectrum. The fraction of SUs that accept the currently advertised price by the provider is dictated by an underlying demand function. The provider must ensure that admission of SUs does not significantly affect quality of service of PUs. This is because presence of SUs may increase blocking of PU calls and hence lead to a punishment in the form of loss of business due to poor service.

The problem of pricing of shared resources has been widely studied in the literature [3–5]. Recent works introduced pricing strategies specifically tailored for secondary access of resources [6–8]. Yet, the overwhelming majority of papers in this area assume that the demand function of users is known (see Sec. II for exceptions). Precise knowledge of the demand function, which may vary over time, is, however, hard to acquire, and raises the question of how to apply this body of existing work in practice.

Furthermore, much of the previous work assume (for analytical tractability) that call lengths are exponentially distributed (see again Sec. II for exceptions). A recent study based on measurement of real traces in a cellular networks shows that this assumption does not hold in practice [9]. In particular, [9] observes that the variance of call length is significantly higher than that of the exponential distribution and questions the validity of previous work (and [6] in particular) based on the exponential distribution assumption.

In this paper, we develop an on-line, measurement-based pricing framework for secondary access applicable to settings where both the demand function and the call length distribution are unknown. We focus on pricing policies, which we refer to as *occupancy-based*, that depend only on the total number of ongoing (SU and PU) calls in the system. We prove that occupancy-based policies are insensitive to the call length distribution, except through the mean. The proof follows through the establishment of a connection with coordinate convex policies in call admission control that are known to enjoy product-form equilibrium distributions and to be insensitive to the call length distribution [10].

This work was supported in part by the US National Science Foundation under grants CNS-0721860, CCF-0916892 and CNS-0238397

Next, we introduce occupancy-based dynamic and threshold pricing policies. In dynamic pricing, a spectrum provider sets a different price for each occupancy level. In threshold pricing, the provider sets a single price and admits an SU only when the occupancy level is below a certain threshold. For the special case of exponentially distributed call lengths, these policies were studied in [6], which shows that optimal threshold pricing performs very close to optimal dynamic pricing. Moreover, the profit region of threshold pricing (i.e., the range of PU arrival rate for which a positive profit can be achieved) is optimal. Due to the insensitivity of occupancy-based policies to the call length distribution, these results directly apply to the general call length distribution case that we study in this paper.

We conduct extensive simulations for call lengths with phase-type distributions exhibiting higher and lower variability than the exponential distribution and various SU demand functions (convex, concave, linear). Through these simulations, we show that the average profits obtained by the optimal occupancy-based dynamic and threshold pricing policies are close to that obtained by the optimal general pricing policy. Note that the optimal general policy is assumed to have exact knowledge of the call length distribution and current phase of each ongoing call, an information that is very difficult to obtain in practice.

Next, we propose a new on-line algorithm, called *Measurement-based Threshold Pricing* (`MTP`), for efficiently pricing secondary spectrum access when the demand function is unknown, but satisfies certain mild assumptions. This algorithm requires the optimization of only two parameters (threshold and price), which makes it highly preferable over trying to mimic an optimal dynamic policy requiring the optimization of a different price for each channel occupancy level.

`MTP` is an iterative algorithm, based on a variation of Fibonacci search, that aims to optimize an unknown profit function which depends on price only. At each iteration, `MTP` measures the average arrival rate of SUs corresponding to a certain *test* price. We show that these price-based measurements are sufficient to derive both the optimal price and the optimal threshold. Though the profit function may be multimodal, we show that `MTP` converges to a local optimum as fast as if the function were unimodal. Specifically, we show that the number of iterations and measurements required by `MTP` are logarithmic in the total number of possible prices and are independent of other variables, such as the total number of channels.

Through simulation, we evaluate the performance of `MTP` with finite measurement windows, which implies that the estimation of the SUs arrival rate at each iteration is noisy. Defining the mean call length to be one unit of time and setting the measurement window length to be one time unit as well, we show that, on average, `MTP` converges to a profit within 10% of optimal threshold pricing within only 5 time units, assuming a range of 10^4 different possible prices. Larger measurement windows of length 10 and 100 time units, bring the average profit of MTP within 4% and 2% of the optimal threshold pricing, respectively.

The rest of this paper is structured as follows. In Section II, we discuss related work. In Section III, we introduce the system model and problem formulation. Our contribution on pricing for generally distributed call lengths is given in Section IV. In Section V, the `MTP` algorithm for pricing with unknown demand function is explained. We provide numerical examples in Section VI and conclude the paper in section VII.

II. Related Work

The related work can be broken down to four main categories: congestion-based pricing, secondary access pricing, unknown demand function and insensitivity to call length distribution.

We start with the well studied area of congestion-based pricing. As such, we restrict our literature review to those papers that are the most relevant. Ref. [3] studies pricing of network resources when arrival rate of all users can be regulated with price. They show that static pricing (a single price is advertised regardless of occupancy level) achieves good performance and is optimal in some asymptotic regimes. This result was extended in [4] in the context of large network asymptotics.

In addition to [6], which we mentioned in the previous section, the following papers consider secondary access pricing. Ref. [7] studies optimal and static pricing policies within the context of a generic rental management optimization problem with two types of customers, which are akin to our SUs and PUs. Ref. [11] provides a game theoretic analysis of revenue maximization problem for secondary spectrum access. Ref. [8] studies secondary spectrum access pricing strategies capturing the effects of network-wide interferences. When applicable, the previous work mentioned above assume a known demand function and exponentially distributed call lengths, on the contrary to the model presented in this paper.

Next, we present related work on less studied field of pricing with unknown demand function. Ref. [12] introduces an on-line algorithm for static pricing of calls with exponentially distributed call lengths. It considers a parametric demand function while we consider a more general non-parametric demand function. Ref. [13] studies a different model than ours where the pricing problem is finite horizon and there is a single product with a finite inventory. They provide on-line learning algorithms for parametric and non-parametric demand functions.

Finally, we provide related work on insensitivity to call length distribution. Ref. [5] shows that static pricing policy is insensitive to call length distribution and is still asymptotically optimal. Ref. [14] also studies static pricing for generally distributed call lengths and shows that the profit function of static pricing is unimodal. It assumes that the demand function is known and that all users are elastic to the price, contrary to our model. Refs. [10] study optimal call admission policy for generally distributed call lengths whereas we study pricing policies. More information on insensitivity to call length distribution can be found in [15, 16].

III. Model and Problem Formulation

In this section, we introduce our model and objective. We consider a single-cell wireless network which provides access to C channels. Calls from PUs arrive according to a Poisson process with fixed rate $\lambda_p > 0$. A punishment in the amount of K monetary units is imposed if all the channels are busy and a PU call is blocked. SU call arrivals also form a Poisson process with rate $\lambda_{SU} > 0$ that is independent of the PUs. We note that the measurement study in [9] justifies the use of the Poisson process to model call arrival rates. When an SU call arrives, it accepts with probability $p(u)$ the price u advertised by the provider and attempts to join the network. Therefore, the rate at which SUs attempt to access the spectrum is $\lambda_s(u) = \lambda_{SU} p(u)$. We refer to $\lambda_s(u)$ as the *demand function*.

Some of the results in this paper assume one or both of the following assumptions on the demand function. We specifically state whenever these assumptions are required.

Assumption 3.1: There exists a maximum price $u_{\max}$ for which $\lambda_s(u_{\max}) = 0$. Moreover, $\lambda_s(u)$ is a strictly decreasing, differentiable function in u over the interval $[0, u_{\max}]$.

The second assumption enables development of our efficient on-line optimization procedure presented in Section V.

Assumption 3.2: Let $u(\lambda_s)$ be the inverse of $\lambda_s(u)$ on the interval $0 \le u \le u_{\max}$. Then $\lambda_s u(\lambda_s)$ is concave with respect to λ_s.

Assumption 3.2 implies that the marginal instantaneous profit is decreasing with respect to user demand. It ensures a well-behaved demand function [17]. This assumption is widely made in the literature [3, 13, 17] and is satisfied by variety of demand functions such as functions with exponential, linear and polynomial decay.

We assume that PU and SU call lengths have a common general distribution with mean $1/\mu$. Therefore, once accepted PU and SU calls are statistically indistinguishable. The call length distribution, except its mean, is unknown. Without loss of generality, we assume that $\mu = 1$, i.e., the mean call length time is one unit of time.

In this paper, we restrict our attention to pricing policies that are based solely on the total number (SU and PU) of ongoing calls in the system. We refer to these policies as *occupancy-based* policies. Note that the total occupancy is not Markovian unless call lengths are exponentially distributed; hence an optimal policy would typically entail further information such as the amount of time each call has already been in the system. An occupancy-based pricing policy sets a advertised SU price u_n when there are $n < C$ ongoing calls in the system. Therefore, a policy can be defined as a vector $\mathbf{u} = (u_0, u_1, u_2, ..., u_{C-1})$. We are interested in finding the vector $\mathbf{u}$ which maximizes the average profit per unit of time gained from accepting SUs. Prices for $\mathbf{u}$ is selected from a discrete set $\mathbb{U}$ taking values in the interval $[0, u_{max}]$.

We limit optimal pricing policy search to occupancy-based policies for practical concerns; the specific form of the call length distribution is often unavailable or cannot be properly formalized. Even in such cases where the distribution is known, it is hard to price optimally due to uncertainty in the future length of ongoing calls. On the other hand, we show that occupancy-based policies are insensitive to call length distribution. In Section VI, we provide numerical examples showing that the optimal occupancy-based policy performs very close to the optimal general policy.

IV. Pricing with general call length distribution

In this section, we show that occupancy-based pricing policies are insensitive to call length distribution except through the mean. We prove this property by showing the insensitivity of the equilibrium probability distribution of occupancy.

A. Insensitivity property

First, we make a connection between our pricing problem and the following general resource sharing problem. Assume there are J classes of calls and C available channels. The state of the system, $\mathbf{n} = (n_1, n_2, ..., n_J)$ where $\sum_{j=1}^{J} n_j \le C$, is defined in terms of the number of ongoing class j calls n_j. When the system is in state $\mathbf{n}$, calls of class j arrive according to a Poisson process with rate $\lambda_j(\mathbf{n})$. Note that the arrival rate is state dependent. Call lengths follow a general distribution with mean $1/\mu_j$. A resource sharing policy is a set of rules which dictates whether or not to accept an incoming call of class j when the system in state $\mathbf{n}$.

We are specifically interested in *coordinate convex policies*, defined below, which have the useful property of a product form equilibrium probability distribution that is insensitive to the call length distribution [10].

Definition 4.1: Let Ω be the set of all possible states $\mathbf{n}$. A policy is coordinate convex if there exist a subset $\omega \subseteq \Omega$ such that 1) $\mathbf{n} \in \omega$ and $n_j > 0$ imply $(n_1, ..., n_j - 1, ..., n_J) \in \omega$, 2) it accepts a call of class j when in state $\mathbf{n}$ if $(n_1, ..., n_j + 1, ..., n_J) \in \omega$ for $j = 1, 2, ..., J$.

Next, we state our main result in this section which is proven by establishing a connection between occupancy-based pricing policies and coordinate convex polices.

Theorem 4.2: For any occupancy-based pricing policy $\mathbf{u}$, the equilibrium probability distribution of occupancy is insensitive to call length distribution except through its mean and has the following product form:

$$\pi_n = \frac{\frac{\lambda(0)\lambda(1)\lambda(2)...\lambda(n-1)}{n!}}{1 + \frac{\lambda(0)}{1!} + \frac{\lambda(0)\lambda(1)}{2!} + \ldots + \frac{\lambda(0)\lambda(1)\lambda(2)...\lambda(C-1)}{C!}}, \quad (1)$$

where $\lambda(n) = \lambda_p + \lambda_s(u_n)$ and n is the number of ongoing calls in the system.

Proof: Suppose $J = 1$, i.e., there is a single class of calls. Consider a policy which always accepts a call unless the system is full. This policy satisfies both conditions in Definition 4.1 because both $n + 1$ and $n - 1$ are allowable states whenever $n + 1 \le C$ and $n - 1 \ge 0$. In our model, PU and SU have the same call length distribution and thus are indistinguishable once in the system. Therefore, SU and PU calls can be viewed as belonging to the same class of calls with arrival rate $\lambda(n) = \lambda_p + \lambda_s(u_n)$ at state n. Then, any fixed policy $\mathbf{u}$ always accepts these single class calls. Therefore, it has a product form equilibrium probability distribution which is insensitive to the call length distribution. ■

B. Optimal occupancy-based dynamic and threshold pricing policies

From Theorem 4.2, we deduce that an occupancy-based pricing policy induces the same equilibrium occupancy for all call length distributions with mean 1. Hence one may as well study the equilibrium under exponential call length distribution. Based on the work in [6] with exponentially distributed call lengths, we next summarize the main properties of the optimal occupancy-based dynamic and threshold pricing policies that are necessary to follow the rest of this paper. The results in this section require Assumption 3.1 on the demand function.

First, we provide the average profit function for a given occupancy-based pricing policy with price vector $\mathbf{u}$:

$$R = \sum_{n=0}^{C-1} \pi_n \lambda_s(u_n) u_n - \pi_C \lambda_p K + E(\lambda_p, C)\lambda_p K, \qquad (2)$$

where $E(\lambda_p, C)$ is the blocking probability of PUs in the absence of SU arrivals. This quantity corresponds to the well-known *Erlang-B* formula

$$E(\lambda_p, C) = \frac{\frac{\lambda_p^C}{C!}}{\sum_{n=0}^{C} \frac{\lambda_p^n}{n!}}. \qquad (3)$$

The first term in Eq. (2) represents the sum of the average revenues collected from SUs in each state. The second term is the average punishment due to rejected PUs. The last term $E(\lambda_p, C)\lambda_p K$ acts as the normalization term to ensure that the profit is zero when all SUs are rejected.

Note that the average profit function in Eq. (2) as well as the optimal occupancy-based policy that maximizes it do not depend on the call length distribution except through the mean. Thus, the optimal occupancy-based policy can be calculated by assuming that the call lengths are exponentially distributed. Under this assumption, occupancy-based pricing can be modeled as an average reward dynamic programming problem with C states, and the optimal prices can be calculated by using policy iteration [6]. Since the optimal occupancy-based policy is insensitive to the call length distribution, it is valid for any distribution.

In a threshold pricing policy, SU calls are admitted and charged a price u when the channel occupancy is smaller than some threshold T and rejected otherwise. This is equivalent to having the following price vector

$$\mathbf{u} = (\underbrace{u, u, ..., u}_{T}, \underbrace{u_{max}, u_{max}, ..., u_{max}}_{C-T}).$$

Consequently, the total arrival rate until the occupancy level reaches T channels is $\lambda_p + \lambda_s(u)$ and λ_p afterwards.

Assuming the above price vector, the average profit function of threshold pricing policy is

$$R_T(u) = \sum_{n=0}^{T-1} \pi_n \lambda_s(u) u - \pi_C \lambda_p K + E(\lambda_p, C)\lambda_p K. \qquad (4)$$

The optimization of threshold pricing policy involves finding the optimal values for the price u and threshold T. Since the profit function of threshold pricing policy is insensitive to the call length distribution, it retains properties that are introduced in [6]. In the next section, we will exploit these properties to maximize $R_T(u)$ when the demand function is unknown.

V. Spectrum pricing with unknown demand function

Typically, the SU demand function $\lambda_s(u)$ is unknown. In this section, we introduce an algorithm called Measurement-based Threshold Pricing (MTP) algorithm to calculate the threshold pricing policy under this condition.

When $\lambda_s(u)$ is unknown, a formula for the threshold pricing profit function $R_T(u)$ is unavailable. However, we can measure the arrival rate of SUs for a specific price u and threshold T and calculate the average profit $R_T(u)$ for that price and threshold. Measurements are conducted by observing, for a sufficiently long period of time, the rate of SUs who accept the advertised price. In this section, we assume that measurements are exact. In Section 6, we numerically study the robustness of MTP to noise due to finite measurement windows. The threshold T used during the measurement is irrelevant due to the following property of $R_T(u)$.

Lemma 5.1: For a given price u, $R_T(u)$ can be calculated for any threshold $1 \leq T \leq C$ with a single measurement.
This lemma is direct consequence of Eq. (4), which can be calculated for any threshold and a given price once the corresponding $\lambda_s(u)$ is acquired as a result of a measurement.

In practice, required measurements have to be performed while the system is in operation. These measurements are often done with non-optimal parameters which causes profit loss. Therefore, our main goal is to calculate the optimal threshold pricing policy with as few measurements as possible.

A. Properties of threshold pricing profit function

The properties we introduce in this section require both Assumption 3.1 and Assumption 3.2 on the demand function. These assumptions ensure that, for a fixed threshold, the profit function $R_T(u)$ is unimodal with respect to price in $[0, u_{max}]$ [6]. A function is unimodal over a certain interval, if it has a single maximum over that interval. This property enables efficient calculation of the optimal price for a given threshold. However, finding the optimal threshold requires a search over all possible threshold values. We circumvent this problem by introducing an auxiliary profit function which depends on price only:

$$R_{max}(u) = \max_{1 \leq T \leq C}(R_T(u)) \qquad (5)$$

For a given price u, this function can be calculated with a single measurement thanks to Lemma 5.1.

During our numerical studies, we observed that, for certain range of system parameters, $R_{max}(u)$ possesses the following property.

Claim 5.2: $R_{max}(u)$ can be multimodal in u for certain system parameters and demand functions.

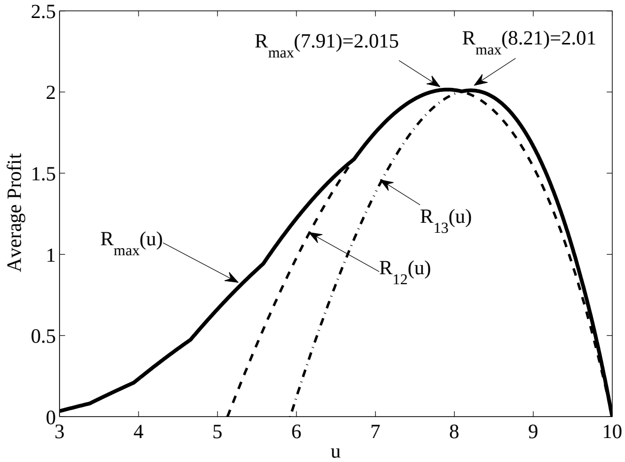

Fig. 1. Multimodal $R_{max}(u)$ and $R_T(u)$ for $T = 12$ and $T = 13$ on which maximums occur. System parameters are $C = 20$, $\lambda_s(u) = (10-u)_+$, $\lambda_p = 12.5$ and $K = 120$.

Proof: Consider a 20 channel system with linear demand function $\lambda_s(u) = (10-u)_+$. The PU arrival rate is $\lambda_p = 12.5$ and the penalty for blocking a PU user is $K = 120$. The function $R_{max}(u)$, for this set up, is plotted in Figure 1. This specific function has two maximum points at $u = 7.91$ and at $u = 8.21$. The corresponding maximizing thresholds are $T = 12$ and $T = 13$, respectively. ■

Claim 5.2 is a rather undesirable property from an optimization point of view. Nevertheless, in the next section, we show that a local optimal price and threshold can be calculated as efficiently as if $R_{max}(u)$ were unimodal.

B. Measurement-based Threshold Pricing (`MTP`*)*

In this section, we describe the `MTP` algorithm and prove that it converges to a local maximum of $R_{max}(u)$. As expected, when the function is unimodal it converges to the global maximum. During our numerical studies, we observed that when $R_{max}(u)$ is multimodal, the average profits of local maximums are very close to each other, as observed in Fig. 1. Therefore, we do not expect significant profit loss when the algorithm converges to a local maximum rather than the global one.

While one can calculate the value of $R_{max}(u)$ for a given price with a single measurement, same is not true for its derivative which can be undefined at certain points (transition points from one $R_T(u)$ to another). Therefore, we base the `MTP` algorithm on the derivative-free *Fibonacci search* which was first introduced by Kiefer [18]. Fibonacci search is a sequential line search algorithm which maximizes a unimodal function. In every iteration, it makes a function evaluation. Together with the information from earlier evaluations, it reduces the minimum interval where the optimal point is known to lie. This interval is referred to as *interval of uncertainty*. Under the following criteria of optimality, Fibonacci search is optimal for searching the maximum of a unimodal function. If the number of function evaluations is fixed in advance, Fibonacci search finishes with the largest ratio of initial size of interval of uncertainty to its final size [18].

In our case, function evaluations, i.e., measurements are conducted for discrete values of price. Therefore, we utilize a discrete version of Fibonacci search (also known as lattice search)[19]. While Fibonacci search might fail to converge when the function is multimodal, `MTP` converges to a local maximum of $R_{max}(u)$. We manage this by taking advantage of the fact that $R_{max}(u)$ is the maximum of unimodal functions $R_T(u)$ for $1 \leq T \leq C$. Algorithm 1 provides a pseudo-code for `MTP` which we next explain.

In the $i^{th} \geq 0$ iteration, `MTP` attempts to maximize the unimodal function $R_{T_i^*}(u)$ the same way as Fibonacci search would do. Here, T_i^* represents the *active threshold* in iteration i which we calculate in the following manner. Let S be the set of prices for which measurements have been obtained so far, i.e., if $u \in S$, then we know the corresponding arrival rate $\lambda_s(u)$. For $u \in S$, we can then calculate $R_T(u)$ for all values of T and deduce the value of $R_{max}(u)$ as well. Let

$$u_i^* = \arg\max_{u \in S}(R_{max}(u)) \tag{6}$$

be the price which yields the maximum profit observed so far. Then,

$$T_i^* = \arg\max_{1 \leq T \leq C}(R_T(u_i^*)) \tag{7}$$

is the optimal threshold for the price u_i^* and $R_{T_i^*}(u_i^*) = R_{max}(u_i^*)$ is the maximum profit calculated so far. In every iteration, `MTP` makes a measurement for a new test price. At the end of every iteration, the active threshold is updated according to this new measurement.

`MTP` chooses the new test price according to Fibonacci numbers. Fibonacci numbers are defined such that

$$F_k = F_{k-1} + F_{k-2}$$

where $F_0 = 0$ and $F_1 = 1$. Let, $\hat{\mathbb{U}}^i$ be the interval of uncertainty in the i^{th} iteration. `MTP` requires that the initial interval of uncertainty $\hat{\mathbb{U}}^0$ contains exactly F_m+1 prices where m is the smallest integer which satisfies $|\mathbb{U}| \leq F_m + 1$. Recall that $\mathbb{U}$ is the set of all possible prices. In order to comply with this condition, we insert $F_m + 1 - |\mathbb{U}|$ fictitious prices to the end of the price series. We assume that the fictitious prices are all equal to u_{max}.

Let u_j^i represent the $j^{th} \geq 0$ price in $\hat{\mathbb{U}}^i$. Then, $\hat{\mathbb{U}}^0 = (u_0^0, u_1^0, u_2^0, .., u_{F_m}^0)$ which naturally contains all local optima of $R_{max}(u)$. In every iteration, the size of the interval of uncertainty is reduced such that $|\hat{\mathbb{U}}^i| = F_{m-i} + 1$ i.e.,

$$\hat{\mathbb{U}}^i = (u_0^i, u_1^i, u_2^i, ..., u_{F_{m-i}}^i) = [u_0^i, u_{F_{m-i}}^i].$$

`MTP` reduces $\hat{\mathbb{U}}^i$ by comparing $R_{T_i^*}(u)$ for two internal test prices, $u_{F_{m-i-2}}^i$ and $u_{F_{m-i-1}}^i$. For the sake of simpler notation, we denote these prices as u_a^i and u_b^i, respectively.

The algorithm starts with an initialization step in which m is calculated and $\hat{\mathbb{U}}^0$ is constructed. Then, measurements for

Algorithm 1 Measurement-based Threshold Pricing (`MTP`)

Calculate m and construct $\hat{\mathbb{U}}_0$
Make measurements for u_a^0 and u_b^0
$u_0^* = \arg\max_{u \in S}(R_{max}(u))$
$T_0^* = \arg\max_{1 \leq T \leq C}(R_T(u_0^*))$
for $i = 0$ to $m-4$ **do**
  if $R_{T_i^*}(u_a^i) \geq R_{T_i^*}(u_b^i)$ **then**
    $\hat{\mathbb{U}}^{i+1} = [u_0^i, u_b^i]$
    Make measurement for u_a^{i+1}
  else
    $\hat{\mathbb{U}}^{i+1} = [u_a^i, u_{F_{m-i}}^i]$
    Make measurement for u_b^{i+1}
  end if
  $u_{i+1}^* = \arg\max_{u \in S}(R_{max}(u))$
  $T_{i+1}^* = \arg\max_{1 \leq T \leq C}(R_T(u_{i+1}^*))$
end for
return u_{m-3}^* and T_{m-3}^*

u_a^0 and u_b^0 are obtained. The initialization step ends with the calculation of u_0^* and T_0^*.

Each iteration starts by comparing the value of $R_{T_i^*}(u_a^i)$ and $R_{T_i^*}(u_b^i)$. If $R_{T_i^*}(u_a^i) \geq R_{T_i^*}(u_b^i)$, then we have

$$R_{T_i^*}(u_a^i) \geq R_{T_i^*}(u_b^i) \geq R_{T_i^*}(u_{F_{m-i}}^i).$$

Since $R_{T_i^*}(u)$ is a unimodal function, the optimal price for $R_{T_i^*}(u)$ can not be in $(u_b^i, u_{F_{m-i}}^i]$. Therefore, the interval of uncertainty is reduced to $\hat{\mathbb{U}}^{i+1} = [u_0^i, u_b^i]$. For the next iteration we need $u_a^{i+1} \in S$ and $u_b^{i+1} \in S$. It can be verified that $u_b^{i+1} = u_a^i$, and thus we only need to make a measurement for u_a^{i+1}.

If $R_{T_i^*}(u_a^i) < R_{T_i^*}(u_b^i)$, we have

$$R_{T_i^*}(u_0^i) \leq R_{T_i^*}(u_a^i) < R_{T_i^*}(u_b^i).$$

Due to similar arguments to those in the previous case, the interval of uncertainty is reduced to $\hat{\mathbb{U}}^{i+1} = [u_a^i, u_{F_{m-i}}^i]$. In this case $u_a^{i+1} = u_b^i$ and we make a measurement for u_b^{i+1}.

At the end of each iteration, u_{i+1}^* and T_{i+1}^* are updated for the use of next iteration. The algorithm terminates after $m-3$ iterations (when $i = m-4$), and returns u_{m-3}^* and T_{m-3}^*, which are local optimal price and threshold of $R_{max}(u)$, as proven next.

We start with the following lemma.

Lemma 5.3: When the active threshold is changed to $T_i^* \neq T_{i-1}^*$, the optimal price for $R_{T_i^*}(u)$ is in $\hat{\mathbb{U}}^i$.

Proof: Assume that the active threshold is changed to T_i^* due to the measurement for price u_a^i. This means, $R_{T_i^*}(u_a^i)$ is the maximum profit calculated so far, i.e., $R_{T_i^*}(u_a^i) > R_{T_i^*}(u_0^i)$ and $R_{T_i^*}(u_a^i) > R_{T_i^*}(u_{F_{m-i}}^i)$ Since $u_0^i < u_a^i < u_{F_{m-i}}^i$ and $R_{T_i^*}(u)$ is unimodal, the optimal price for $R_{T_i^*}(u)$ must be in $\hat{\mathbb{U}}^i$. Same arguments are true if the measurement had been conducted for u_b^i. ■

Next, we present the convergence theorem.

Theorem 5.4: `MTP` converges to a local optimal of $R_{max}(u)$ in $m-3$ iterations and requires $m-1$ measurements, where

$$m = \min_k \{k : |\mathbb{U}| \leq F_k + 1\}.$$

Proof: We first prove the first part of the theorem. In the last iteration $i = m-4$, the interval of uncertainty is reduced to $\hat{\mathbb{U}}^{m-3}$ which contains $|\hat{\mathbb{U}}^{m-3}| = 4$ different prices, and a measurement is conducted for the only price in $\hat{\mathbb{U}}^{m-3}$ which has not been yet tested (either u_a^{m-3} or u_b^{m-3}). Finally, u_{m-3}^* and T_{m-3}^* are calculated. Even though T_{m-3}^* could be different from T_{m-4}^* (the last active threshold), u_{m-3}^* is the optimal price of $R_{T_{m-3}^*}(u)$ due to Lemma 5.3 and the fact that u_{m-3}^* is the best performing price in $\hat{\mathbb{U}}^{m-3}$. u_{m-3}^* is a local optimal of $R_{max}(u)$ because it is the optimal price of $R_{T_{m-3}^*}(u)$ and T_{m-3}^* is the optimal threshold for u_{m-3}^*.

As for the second part of the theorem, `MTP` makes a new measurement in every iteration. Together with the initial two, the algorithm requires $m-1$ measurements. ■

In conclusion to this section, we note that the number of measurements required by `MTP` is the same as in the Fibonacci search which is $\log_\phi(|\mathbb{U}|)+O(1)$ where $\phi = (1+\sqrt{5})/2$ is the golden ratio [20]. `MTP` can easily be adapted to converge to the global maximum of $R_{max}(u)$. To do so, the threshold should be kept fixed throughout the `MTP` algorithm. This should be repeated for all possible thresholds $1 \leq T \leq C$. However, in this case the algorithm would require $C(\log_\phi(|\mathbb{U}|) + O(1))$ measurements instead.

VI. Numerical results

A. Comparison between the optimal occupancy-based and the optimal general policies

In this section, we compare the optimal occupancy-based, the optimal threshold and the optimal general pricing policies. We demonstrate near-optimal performances of occupancy-based policies for diverse set of call length distributions and demand functions.

We consider phase-type call length distributions with two phases, namely, hyper-exponential and hypo-exponential distributions. These distributions provide valuable insight on the performance of occupancy-based policies for broad range of call length variances. For exponentially distributed call lengths, for which the optimal occupancy-based policy is also the optimal general policy, the coefficient of variation c_v (ratio of standard deviation to mean) is 1. These distributions represent divergence from exponential distribution in both directions. For hyper-exponential $c_v \geq 1$ and for hypo-exponential $c_v < 1$ [21].

If a call length has a two-phase hyper-exponential distribution then with probability p it is exponentially distributed with mean $1/\mu_1$ and with probability $1-p$ it is exponentially distributed with mean $1/\mu_2$. In the hypo-exponential case (also known as generalized Erlang distribution), a call starts the first phase which is exponentially distributed with mean $1/\mu_1$, and then continues to the second phase which is exponentially distributed with mean $1/\mu_2$.

TABLE I
PHASE-TYPE DISTRIBUTIONS

Distribution	μ_1	μ_2	p	Mean	Variance
Hyper-1	3	1/3	3/4	1	11/3
Hyper-2	2	2/3	1/2	1	3/2
Hypo-1	2	2	n/a	1	1/2
Hypo-2	10/9	10	n/a	1	41/50

In Table I, we present parameters of the specific phase-type distributions that we consider. We calculate the optimal general policy for these distributions by assuming that we have exact knowledge of the distribution as well the phase of each ongoing call, an information which is hard to obtain in practice. With this information, we model the system as an average reward dynamic programming problem where states are defined as the total number ongoing calls in each phase. It can be shown that there are $(C+1)(C+2)/2$ states and $C(C+1)/2$ prices to optimize, substantially larger than the C states and $C-1$ prices that the occupancy-based model requires.

To demonstrate performances of the pricing policies for SU demand functions with different characteristics, we use the following function which can take various shapes depending on the parameter β:

$$\lambda_s(u) = \alpha\big(\frac{u_{max} - u}{u_{max}}\big)^{\beta}, \qquad u \in [0, u_{max}]. \tag{8}$$

For $\beta = 1$ it is linear, for $\beta < 1$ it is convex and for $\beta > 1$ it is concave. This demand function is also used in [7] and satisfies Assumptions 3.1 and 3.2. We study different shapes of this demand function for $\alpha = 10$ and $u_{max} = 10$.

In Fig. 2, we compare the performances of the optimal occupancy-based, the optimal threshold and the optimal general policy for a linear demand function $(10-u)_+$ and call lengths with phase-type distributions shown in Table I. Fig. 2 shows the average profit generated for different values of λ_p. It is clear that in all the cases, the curves are very close. The greatest difference between the optimal occupancy-based policy and the optimal general policy is observed for the case of Hyper-1, which has the highest variance. To quantify the difference, we compute, for each policy, the average profit taken over all PU arrival rates λ_p for which a strictly positive profit is obtained. In other words, we compute the integral of each function divided by the support of the function (this support happens to be the same for all policies). For the Hyper-1 case, this quantity turns out to be 13.977 (monetary units) for the optimal general policy, 13.836 for the optimal occupancy-based policy, and 13.672 for the optimal threshold policy, i.e., a revenue loss of only about 1% and 2.2% for the latter two policies, respectively.

In [9], Willkomm et. al. conclude that the call length distribution can not be properly modeled by using exponential or Erlang distributions (a special case of hypo-exponential distributions with $\mu_1 = \mu_2$). Moreover, they observed that the variance is more that three times the mean. Therefore, the Hyper-1 model which has variance 3.67 times the mean is consistent with these observations.

We have also conducted numerical analysis for call lengths having two-phase hyper-exponential distribution with higher variances. We observed that, for these higher variances, the performance of the optimal occupancy-based policy remains similar to that in Fig. 2.

In Fig. 3, we compare the performances of the aforementioned policies for concave ($\lambda_s(u) = 10(\frac{10-u}{10})^2$) and convex ($\lambda_s(u) = 10(\frac{10-u}{10})^{1/2}$) demand functions, using Hyper-1 as the call length distribution. For this setting, we observe that the optimal occupancy-based policy as well as the optimal threshold policy again perform close to optimal general policy, the concave case being the worse among the two. In that case, the average profit taken over profit-yielding values of λ_p is 8.163 for optimal general policy, while that of optimal occupancy-based policy is 8.071 (about 1.1% smaller) and that of the optimal threshold policy is 7.970 (2.3% smaller).

B. *Pricing with unknown demand function*

In this section we evaluate the performance of `MTP` with finite measurement windows. We consider the interval $\mathbb{U} = [0, 10]$ as the set of available prices where discrete prices are $\Delta u = 10^{-3}$ monetary units apart. Therefore, the total number of possible prices is $|\mathbb{U}| = 10001$. It can be verified that, in this case, $m = 21$ because $F_{21} = 10946$. It means that `MTP` terminates in $m - 3 = 18$ iterations and makes $m - 1 = 20$ measurements.

In Figures 4 and 5, we show the average profit ($R_T(u)$) corresponding to the current test price and the current active threshold T_i^*. We run the algorithm 100 times and take sample average of the average profits. We also display 95% confidence interval for these samples for each measurement window. Recall that, we define the mean call length to be one time unit, which is typically in the order of a few minutes [9].

In Fig. 4, we assume that SUs request access to the spectrum with a rate of $\lambda_s(u) = 10(\frac{10-u}{10})^{1/2}$ calls per time unit, which is unknown to `MTP`. PUs arrive with a known constant rate of $\lambda_p = 8$ calls per time unit. We use a 10 time units measurement window. The figure shows that `MTP` converges within 4% of the optimal threshold policy and within 5% of the optimal occupancy-based policy in 5 iterations. Note that, at some iterations, `MTP` selects test prices such that the average profit drops, as seen in the fourth iteration of Fig. 4.

In Fig. 5, we run `MTP` for the linear demand function $\lambda_s(u) = (10-u)_+$ and different measurement windows (note the logarithmic scale of the x-axis). We observe that for 1 time unit measurement window, `MTP` comes within 10% neighborhood of optimal pricing in just 5 time units. Naturally, the longer the measurement window, the closer `MTP` comes to the optimal threshold policy.

VII. CONCLUSION

In this paper, we investigated pricing of secondary spectrum access with unknown demand function and generally distributed call lengths. First, we introduced occupancy-based

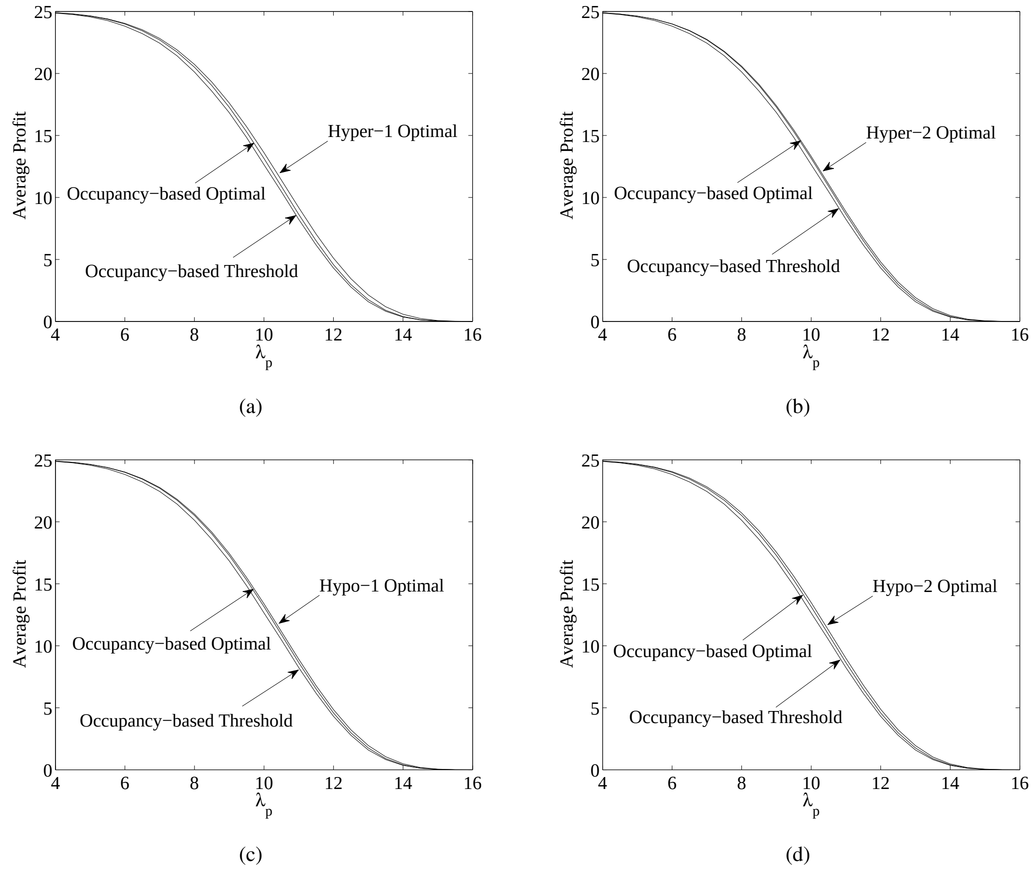

Fig. 2. Comparison of the optimal occupancy-based, the optimal threshold and the optimal general policy for hyper-exponentially and hypo-exponentially distributed call lengths. System parameters: $K = 100$, $C = 20$ and the demand function is $\lambda_s(u) = (u - 10)_+$.

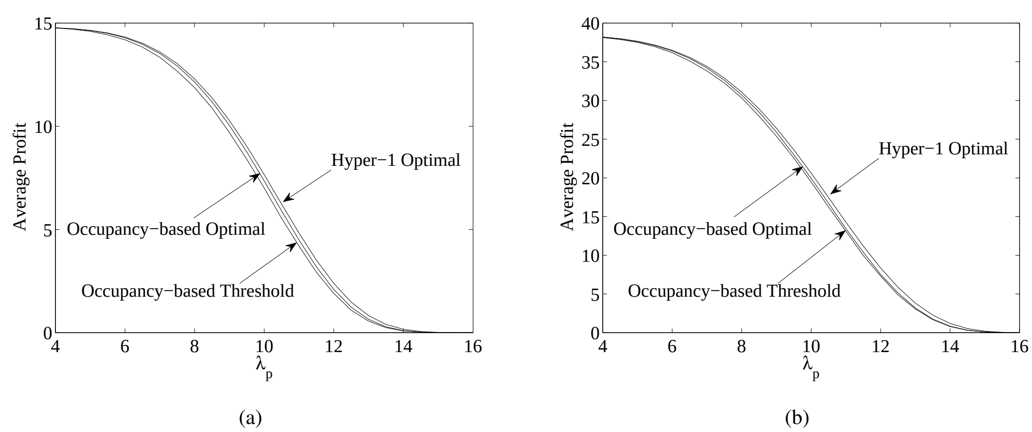

Fig. 3. Comparison of the optimal occupancy-based, the optimal threshold and the optimal general policy for hyper-exponentially distributed call lengths. System parameters: $K = 100$, $C = 20$. In Fig. (a) the demand function is $\lambda_s(u) = 10(\frac{10-u}{10})^2$ and in Fig. (b), it is $\lambda_s(u) = 10(\frac{10-u}{10})^{1/2}$.

pricing policies which price SUs based solely on the total number of ongoing calls. We showed that these policies are insensitive to the call length distribution except through the mean. This property makes occupancy-based policies the solution of choice because the call length distribution is usually unknown or difficult to precisely characterize. We provided numerical comparisons of the optimal occupancy-based, the optimal threshold and the optimal general policy when call lengths have hyper-exponential and hypo-exponential distributions. For these phase-type distributions, we calculated the optimal general policy assuming that we know the exact phase of each ongoing call. We observed that for the optimal

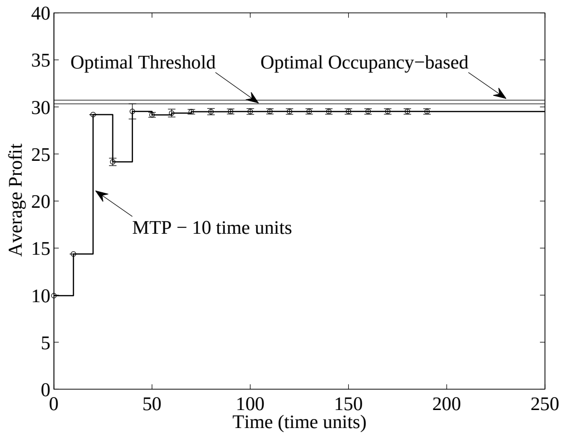

Fig. 4. Performance of MTP with measurement window 10 time units, $\lambda_s(u) = 10(\frac{10-u}{10})^{1/2}$ and $\lambda_p = 8$ calls per time unit, $K = 100$, $C = 20$. Average of 100 runs with 95% confidence interval. $\Delta u = 10^{-3}$, $\mathbb{U} = [0, 10]$ and $|\mathbb{U}| = 10001$

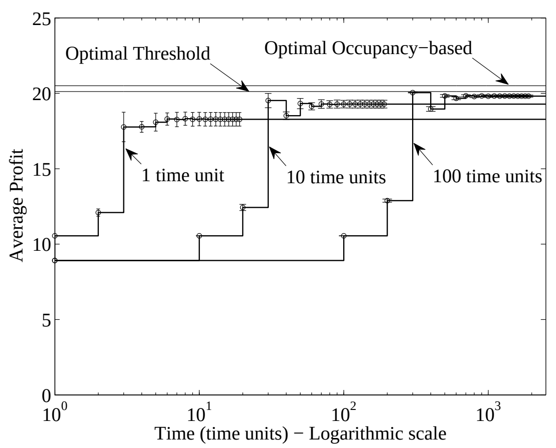

Fig. 5. Performance of MTP with measurement windows 1,10 and 100 time units. $\lambda_s(u) = (10 - u)_+$ and $\lambda_p = 8$ calls per time unit, $K = 100$, $C = 20$. Average of 100 runs with 95% confidence interval. $\Delta u = 10^{-3}$, $\mathbb{U} = [0, 10]$ and $|\mathbb{U}| = 10001$

occupancy-based policy, the average profit over the profit region is within 1% of the optimal general policy, while for the optimal occupancy-based threshold policy it is within 2%.

Second, we devised a new on-line algorithm MTP which converges to a local maximum of $R_{max}(u)$, the threshold policy profit function which depends on price only. In our numerical studies, we observed that $R_{max}(u)$ is generally unimodal and whenever it is multimodal, the local maximum profits are very close to each other. Under mild assumptions on the the demand function, we proved that despite multimodality of $R_{max}(u)$, MTP converges with the same number of measurements as if $R_{max}(u)$ were unimodal which is $\log_\phi(|\mathbb{U}|) + O(1)$. For example, when $|\mathbb{U}| = 10001$, MTP requires only 20 measurements. However, in our numerical studies, we observed that in just 5 measurements MTP comes very close to its final convergence point. The performance of MTP depends on the length of the measurement window. In our numerical result, we observed that with 10 time units measurement windows, in just 5 measurements, MTP comes within 4% of the average profit of the optimal threshold policy which knows the demand function.

In summary, this work provides a novel robustness perspective for practical methods to manage secondary spectrum access. Demonstrated performance of these methods under minimal assumptions on the underlying environment facilitates spectrum reforms in achieving their full potential.

REFERENCES

[1] "The development of secondary markets." http://www.fcc.gov/sptf/headlines2003.html, October 7 2003. Report & Order.

[2] "Meet the eBay of wireless spectrum. Spectrum Bridge automates, webifies wireless deals." http://www.networkworld.com/news/2008/091208-wireless-spectrum.html, September 12 2008. News Release.

[3] I. C. Paschalidis and J. N. Tsitsiklis, "Congestion-dependent pricing of network services," *IEEE/ACM Transactions on Networking*, vol. 8, pp. 171–184, April 2000.

[4] I. C. Paschalidis and Y. Liu, "Pricing in multiservice loss networks: static pricing, asymptotic optimality, and demand substitution effects," *IEEE/ACM Transactions on Networking*, vol. 10, pp. 425–438, June 2002.

[5] X. Lin and N. B. Shroff, "Simplification of network dynamics in large systems," *IEEE/ACM Trans. Netw.*, vol. 13, no. 4, pp. 813–826, 2005.

[6] H. Mutlu, M. Alanyali, and D. Starobinski, "Spot pricing of secondary spectrum access in wireless cellular networks," *IEEE/AMC Transactions on Networking*, To appear, 2009.

[7] N. Gans and S. Savin, "Pricing and capacity rationing for rentals with uncertain durations," *Management Science*, vol. 53, March 2007.

[8] A. Al-Daoud, M. Alanyali, and D. Starobinski, "Secondary pricing of spectrum in cellular CDMA networks," in *IEEE DySPAN*, 2007.

[9] J. B. D. Willkomm, S. Machiraju and A. Wolisz, "Primary users in cellular networks: A large-scale measurement study," in *New Frontier in Dynamic Spectrum Access Networks(DySPAN)*, 2008.

[10] K. W. Ross and D. H. Tsang, "The stochastic knapsack problem," *IEEE Transactions on Communication*, vol. 37, no. 7, pp. 740–747, 1989.

[11] A. O. Ercan, J. Lee, S. Pollin, and J. M. Rabaey, "A revenue enhancing stackelberg game for owners in opportunistic spectrum access," in *IEEE DySPAN*, 2008.

[12] E. Campos-Nanez and S. Patek, "On-line tuning of prices for network services," in *INFOCOM*, 2003.

[13] O. Besbes and A. Zeevi, "Dynamic pricing without knowing the demand function: Risk bounds and near optimal algorithms," *Operations Research*, To appear, 2008.

[14] S. Ziya, H. Ayhan, and R. Foley, "A Note on optimal pricing for finite capacity queueing systems with multiple customer classes," *Naval Research Logistics*, vol. 55, pp. 412–418, 2008.

[15] D. Burman, J. Lehoczky, and Y. Lim, "Insensitivity of blocking probabilites in a circuit-switching network," *Adv. Appl. Prob.*, vol. 21, pp. 850–859, 1984.

[16] J. Kaufman, "Blocking in a shared resource environment," *IEEE Trans. Comm.*, vol. 29, pp. 1474–1481, 1981.

[17] S. Ziya, H. Ayhan, and R. Foley, "Relationships among three assumptions in revenue management," *Operations Research*, vol. 52, pp. 804–809, September-October 2004.

[18] J. Kiefer, "Sequential minimax search for a maximum," *Proceedings of the American Mathematical Society*, vol. 4, no. 3, pp. 502–506, 1953.

[19] D. Wilde, *Optimum seeking methods*. Prentice-Hall, 1964.

[20] E. D. Demaine and S. Langerman, "Optimizing a 2d function satisfying unimodality properties," *Lecture Notes in Computer Science, Algorithms ESA*, vol. 3669, pp. 887–898, 2005.

[21] G. Bolch, S. Greiner, H. deMeer, and K. Trivedi, *Queueing networks and Markov chains*. 2006.