Profitability of Dynamic Spectrum Provision for Secondary Use

Murat Alanyali
Boston University

Ashraf Al Daoud
German-Jordanian University

David Starobinski
Boston University

***Abstract*—We characterize policies and prices for secondary spectrum provision whose profitability is *insensitive* to the demand curve. In more explicit terms, the paper provides *a critical price value* such that if secondary access is priced above that value then allowing secondary access is profitable for the licensee as long as the price generates secondary demand. Conversely, if the price does not generate demand then the licensee does not incur any operational cost due to secondary service. Hence such characterization serves as a guarantee that a spectrum licensee can strictly avoid revenue loss due to participation in spectrum trading.**

I. Introduction

Secondary spectrum markets bear tremendous potential to increase spectrum utilization by allowing spectrum ownership to float in response to varying demand and supply conditions. Realizing this potential entails liquidity of spectrum markets, which in turn requires their profitability for parties that are endowed with the initial allocation of spectrum. Profitability relies on multiple factors, first and foremost on price-demand relationships that are difficult to predict in emerging markets. This paper investigates conditions under which secondary-market transactions are profitable for a spectrum licensee, and gives a constructive perspective for spectrum pricing to achieve profitability.

From the viewpoint of a spectrum licensee, opening a frequency band to secondary access entails determining a pricing strategy and a policy under which such access will be allowed. This pair does not single-handedly determine the net gain from spectrum sharing; the resulting profit also depends on the relationship between the stipulated price of secondary access and the demand it generates. This relationship is coined as the *demand curve*. While the demand is typically non-increasing in price, an exact functional relationship is seldom available without exhaustive measurements, and it may further be time-varying due to market conditions.

The illustration of Fig. I helps clarify the generic issue associated with unknown demand curve. The horizontal and vertical axes of the figure respectively represent the stipulated prices and corresponding demand for

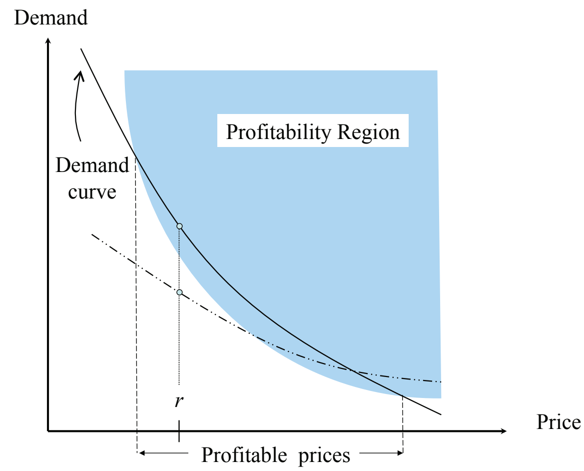

Fig. 1. Graphical illustration of profitability for a generic spectrum sharing policy: Price r yields positive profit under the solid demand curve, but not under the dashed demand curve.

secondary spectrum use. Not all price-demand pairs necessarily guarantee profit since occupancy of spectrum by secondary users may potentially prevent revenue from primary use. Thus, for a given secondary access policy, only a subset of secondary price-demand pairs generate positive profit for the licensee. We refer to this subset as the *profitability region* of secondary spectrum use.

At the boundary of the profitability region the revenue generated from secondary access balances the opportunity cost due to forfeited access rights; in turn the profit is zero. The net effect is typically revenue loss in the remaining portion of the price-demand plane. If the demand curve intersects the profitability region then the policy yields profit provided that secondary access is priced in a way to render the operating point in that region. A range of such prices is illustrated in Fig. I. It is evident that this range generally depends on the demand curve. For example, for the hypothetical policy of this figure the price r is profitable under the solid demand curve, but not under the dashed demand curve.

Our goal in this paper is to characterize policies and prices for secondary spectrum provision whose prof-

This work was funded in part by NSF through grant CCF-0964652.

itability is *insensitive* to the demand curve. In more explicit terms, the paper provides *a critical price value* such that if secondary access is priced above that value then allowing secondary access is profitable for the licensee as long as the price generates secondary demand. Conversely, if the price does not generate demand then the licensee does not incur any operational cost due to secondary service. The critical price obtained here is the smallest price value with this property.

Existence and characterization of a critical price serves as a guarantee that a spectrum licensee can strictly avoid any revenue loss due to participation in spectrum trading. The decoupling between profitability and demand has also implications beyond the basic setting considered in the paper. These include: (i) *Competitive environments:* The critical price is an absolute limit on how low a licensee can price its spectrum; hence it is an indicator of competitiveness of the licensee, particularly in price-war scenarios. (ii) *Service level agreements:* The licensee can continue to honor any service level agreements with the primary demand by first pricing secondary access for profitability and then by rejecting part of the resulting secondary demand as needed. In effect this amounts to reshaping the demand curve and has no bearing on profitability. (iii) *Revenue enhancement:* The licensee can enhance its revenue via multiple rounds of secondary spectrum offerings, at each round pricing above the critical price that is determined by the demand raised in prior rounds. These extensions are discussed in further detail in Section VIII.

A. Overview

As an asset radio spectrum has distinct features that are shaped by physical, technological, and regulatory factors. In this paper we study a model that incorporates the following aspects:

1) *Spatial dimension:* Spectrum is a spatially distributed asset that lends itself to sparse reuse. The sparsity constraint arises due to interference, which depends on a combination of factors including transmission power, frequency, bandwidth, and modulation.
2) *Coexistence:* Multiple types of users can coexist on the same frequency band as long as they comply with operational constraints associated with interference. For example, a cellular network provider may serve its primary customer base while offering temporal surplus of its capacity to secondary users on an opportunistic basis. For another example, multiple users can coexist under the private-commons model [1, Section IV.B.2.b] without using infrastructure of a licensee.
3) *Dynamic transactions:* Currently spectrum transactions in the US are subject to FCC approval that may take weeks. A considerable fraction of these transactions represent transfer of rights for a duration shorter than the remaining lifetime of the original license [2]. Hence for the involved licensees these are dynamic transactions involving the same spectrum block. FCC's vision extends far beyond the present setup and includes per-session contracts that take place at faster time-scales [1, Section IV.B.2.a].

We consider spatial properties of spectrum in terms of a graph where nodes represent geographical locations and edges represent proximity. It is assumed that the spectrum cannot be in use in two neighboring locations simultaneously. On this model we consider a dynamic demand for spectrum that is composed of spectrum requests that occur randomly in time at each location. Each request is for a finite duration which may be different from request to request.

We refer to the spectrum demand prior to a pricing decision as *primary demand*. This demand is subject to a certain primary price, which is the revenue obtained per primary request. It is assumed that primary requests are admitted whenever they do not violate any of the mentioned interference constraints in G. The alluded pricing decision concerns opening spectrum to additional demand at a secondary price. We refer to this demand as *secondary demand*. While the secondary price is at the discretion of the licensee, the licensee has no control on the generated secondary demand. Furthermore the licensee does not know the secondary demand that a given price would raise. Too low a price may generate high secondary demand and consequently may lead to revenue loss due to lost opportunities to serve primary demand; whereas too high a price may not generate any secondary demand at all.

We provide an explicit characterization of a critical price value r^* such that if secondary access is priced above r^* then allowing secondary access is profitable for the licensee as long as the price generates secondary demand. Profitability of such a price entails rather nontrivial secondary access policies that impose further conditions on when to admit secondary requests. We give an explicit characterization of these policies. The profitability threshold r^* is sharp: Any price below this threshold incurs revenue loss if it generates demand, regardless of the adopted spectrum sharing policy. The corresponding profitability region therefore represents an area delineated by a straight vertical line in the secondary price-demand plane.

While the obtained policies yield the largest possible profitability region, neither computational nor operational requirements of these policies scale well with the size of the spatial representation. Hence practical relevance of the optimal profitability region appears limited to smaller topologies. For larger topologies, we obtain the profitability region of a complete sharing policy that imposes no additional admission conditions on the secondary demand. We provide a critical price value r^*_{CS} such that if secondary access is priced above r^*_{CS} then complete sharing is profitable as long as the price generates secondary demand. r^*_{CS} is the smallest such threshold for complete sharing; hence it would raise positive demand if any larger threshold does. The profitability region of complete sharing is not delineated by a straight line in general; therefore profitability of prices below r^*_{CS} depends on the underlying demand curve. We also illustrate how multiple spectrum offerings can be employed to increase the revenue further, and briefly discuss how the service quality of primary traffic can be protected against high secondary demand.

The paper is structured as follows: The next section gives an overview of related work. Sections III and IV specify the adopted spectrum model and profitability definitions respectively. An analysis of spectrum revenue under exclusively primary demand is given in Section V. This revenue serves as a reference value to assess profitability of secondary access. The two profitability regions discussed above are established in Sections VI and VII respectively. Discussions of extensions and computational considerations are included in Section VIII and in an appendix. The paper concludes with final remarks in Section IX.

II. Related work

In contrast to profit maximization, profitability has received scant attention both in economics and in engineering. Existing work on secondary spectrum markets predominantly centers on auction mechanisms and pricing under provision of a centralized broker [3–7]. A notable exception is [8], which studies conditions for market liquidity in exchange-based spectrum trading. The authors of the present paper are not aware of a fundamental analysis of profitability.

Pricing with unknown demand function (curve), with the goal of maximizing profit, has received more attention. Ref. [9] devises an on-line method for optimizing a static (fixed) pricing policy. It assumes a parameterized demand function and adjusts prices in response to customer arrivals. Ref. [10] presents a measurement-based strategy, called Measurement-based Threshold Pricing (MTP), for pricing secondary spectrum access. The paper proves that MTP quickly and provably converges to optimal profit, while making minimal assumptions about the structure of the demand function. However, the MTP framework applies to an isolated cell, and therefore does not capture the spatio-temporal nature of spectrum use and the interference constraints that it induces in general network topologies. A key feature of the model introduced in our paper is to explicitly capture the dynamic fluctuations of spectrum usage in both time and space.

Ref. [11] provides on-line pricing algorithms for both parametric and non-parametric unknown demand functions. The results apply to a different set-up that that considered in our paper, namely finite horizon pricing of a single product with a finite initial inventory. Ref. [12] studies a similar problem, but with competition between multiple providers. Assuming a linear demand function, their approach is based on obtaining the least square estimates of the parameters that define this function.

Our paper stems from extensive work conducted by the telecommunication policy community in identifying and articulating challenges faced towards the deployment of secondary spectrum markets. In a nutshell, [13, 14] detail the multiple dimensions of radio spectrum (e.g., space, time, frequency) that need to be accounted for, [15] discusses market dynamics and approaches for generating demand, and [16–18] describe pros and cons of various spectrum sharing mechanisms (e.g., property rights and commons). We refer to [19, 20] for a comprehensive review of related work in this area.

III. A spatio-temporal model of spectrum

In this paper we represent spatial properties of spectrum sharing via an undirected graph G with n nodes. In this graphical model each vertex represents a geographical location. An edge is drawn between two vertices if the same band of spectrum cannot be utilized simultaneously at the two locations represented by the vertices, due to interference. In that case we call the two vertices (or the two locations) *neighbors* of each other.

In a cellular service context each node of G may represent a cell with an associated base station. Fig. 2 illustrates a 32-node hexagonal lattice topology. Alternatively, in a private commons context, G may represent a suitable discretization of space. In that case the number of edges per node in G would be increasing with the resolution of the spatial representation. The analytical discussion of the paper does not make any assumptions on the topology of G.

We consider the following dynamic demand model superimposed on G: At each location there are two types of spectrum requests. *Primary requests* arrive according to a Poisson process of rate $\lambda_1 > 0$ and *secondary requests* arrive according to a Poisson process of rate $\lambda_2 > 0$. We shall assume that these arrival processes are independent; furthermore that they are independent across different

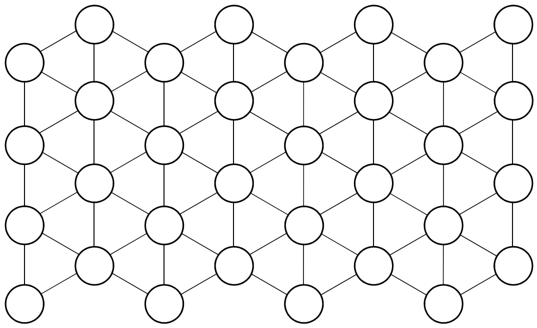

Fig. 2. An example of graph G that represents locations and interference constraints as explained in the text.

locations. If a request at some location i is granted then it holds spectrum for a random duration, during which spectrum is not available to incoming requests at i as well as at neighboring locations of i. We assume that the holding times of requests are independent and identically distributed exponential random variables; and without loss of generality we take the mean holding time as one unit.

Let us denote the occupancy of each location i by x_i, where

$$x_i = \begin{cases} 1 & \text{if there is an active spectrum user} \\ & \text{(primary or secondary) at location } i \\ 0 & \text{else,} \end{cases}$$

and the entire network occupancy by the binary vector $\boldsymbol{x}$, where

$$\boldsymbol{x} = (x_1, x_2, \cdots, x_n).$$

Since neighboring locations cannot be active simultaneously, any feasible network occupancy vector needs to satisfy $x_i x_j = 0$ for neighbors i, j in G. In the parlance of graph theory such a vector is coined an *independent set* of graph G. We refer to the collection of all independent sets of G by S so that necessarily $\boldsymbol{x} \in S$.

For the spectrum provider a granted request generates revenue $r_1 > 0$ if it is primary, and revenue $r_2 > 0$ if it is secondary. These values may reflect deterministic charges per granted request or average charges if, for example, a granted request is charged based on usage.

We identify primary requests with legacy users of the spectrum provider and assume that primary requests are granted whenever possible. That is, a primary request at location i is granted if at the time of the request $x_i = 0$ and $x_j = 0$ for all neighboring locations j of i. Secondary requests, on the other hand, represent opportunity to increase revenue beyond what can be obtained from the primary requests. Towards that end secondary requests may possibly be admitted selectively according to a *spectrum sharing policy*.

IV. Profitability definitions

The objective of this paper is to identify secondary price-demand pairs (r_2, λ_2) under which sharing the spectrum with secondary requests is profitable for the spectrum provider. Profitability of a spectrum sharing policy is defined relative to the spectrum provider's long-term revenue in the absence of any secondary requests. This reference value is also the revenue of a *lock-out policy* that flatly rejects all secondary requests.

Since long-term revenues are likely to increase without bound with increasing term duration, it is more convenient to work with the rate of revenue generation per unit time. We shall denote by R_{LO} the rate at which the lock-out policy generates revenue in the long-term. The parameters that determine R_{LO} are therefore the network graph G, primary demand λ_1, and primary price r_1. Similarly, the revenue rate of a given spectrum sharing policy SP is denoted by $R_{SP}(r_2, \lambda_2)$. Note that this value also depends on G, λ_1, r_1 but this dependence is suppressed for notational convenience.

For a given context G, λ_1, r_1, we shall say that policy SP is *profitable* under secondary price-demand pair (r_2, λ_2) if

$$R_{SP}(r_2, \lambda_2) > R_{LO}.$$

The *profitability region of policy SP* is the set of pairs (r_2, λ_2) for which SP is profitable. Finally, the *full profitability region* (of spectrum sharing) is the union of profitability regions of all possible spectrum sharing policies. Hence a pair (r_2, λ_2) is in the profitability region if and only if it is in the profitability region of some spectrum sharing policy.

V. Lock-out policy

The revenue rate of the lock-out policy, R_{LO}, is a reference value in determining profitability of any spectrum sharing policy. We start with identifying this quantity in terms of the context parameters. As each granted primary request yields revenue r_1, determining R_{LO} hinges on the rate of granted primary requests. In this section we express this latter value in terms of the equilibrium distribution of the network occupancy $\boldsymbol{x}$.

Under the lock-out policy the network occupancy is a Markov process. State transitions of this process are governed by the so-called generator matrix $Q = [Q_{\boldsymbol{x}\boldsymbol{x}'}]_{S\times S}$, which has one entry for every pair of network states $\boldsymbol{x}, \boldsymbol{x}' \in S$. To fully identify entries of Q let $\boldsymbol{e}_i$ denote the binary vector in S where the only non-zero entry is at location i. Then off-diagonal entries of Q are given by

$$Q_{\boldsymbol{x}\boldsymbol{x}'} = \begin{cases} \lambda_1 & \text{if } \boldsymbol{x}' = \boldsymbol{x} + \boldsymbol{e}_i \text{ for some } i \\ 1 & \text{if } \boldsymbol{x}' = \boldsymbol{x} - \boldsymbol{e}_i \text{ for some } i \\ 0 & \text{else,} \end{cases} \quad (1)$$

and diagonal entries of Q are chosen so as to make the row-sums zero.

A probability vector $\pi = \{\pi(\boldsymbol{x}) : \boldsymbol{x} \in S\}$ is an equilibrium distribution for the network occupancy if $\pi Q = 0$. It can be verified by substitution that this equality is satisfied if

$$\pi(\boldsymbol{x}) \;=\; \pi_{\lambda_1}(\boldsymbol{x}) \;=\; \frac{1}{C}\lambda_1^{\sum_i x_i}, \quad (2)$$

for any constant C. The right value of this constant is $C = \sum_{\boldsymbol{x}\in S} \lambda_1^{\sum_i x_i}$, in which case the entries of π sum to 1 and so π_{λ_1} is indeed a probability vector. We point out that our choice of annotating the equilibrium distribution with the parameter λ_1 is motivated by notational convenience in the ensuing discussion.

Let A_i be the probability that a given location i can grant an incoming primary request in equilibrium. Since primary requests are granted whenever possible, a new request at location i is granted if and only if the network state is in the set

$$S_i^+ = \{\boldsymbol{x} : x_i = 0,\ x_j = 0 \text{ for all neighbors } j \text{ of } i\}.$$

Therefore,

$$A_i = \sum_{\boldsymbol{x}\in S_i^+} \pi_{\lambda_1}(\boldsymbol{x}).$$

Since primary requests arrive at rate λ_1 at location i and they are granted with probability A_i, the rate of granted requests at this location is $\lambda_1 A_i$. Therefore, the location generates revenue at rate $r_1\lambda_1 A_i$ per unit time. R_{LO} is the total rate of revenue generation by all locations in the network; hence

$$R_{LO} = r_1\lambda_1 \sum_i A_i.$$

An alternative representation of R_{LO} is obtained via Little's law, which asserts that $\lambda_1 A_i = E_{\lambda_1}[x_i]$. Here E_{λ_1} denotes expectation with respect to the distribution π_{λ_1}. Substituting this equality in the characterization above yields

$$R_{LO} = r_1 E_{\lambda_1}[T]$$

where

$$T = \sum_i x_i$$

is the total network occupancy. We shall find this latter representation more useful from a computational perspective. We postpone the discussion computational complexities to the Appendix and continue in the next section with characterization of the full profitability region.

VI. Full profitability region

Identifying the profitability region by enumerating all possible spectrum sharing policies and by taking the union of their profitability regions is evidently intractable. Here we use the so-called policy improvement procedure of dynamic programming as a succinct alternative.

Policy improvement is a generic iterative technique to solve dynamic optimization problems. It is based on an arbitrary reference policy with known performance, and it seeks an optimal decision to be applied to the first event at each state, provided that the reference policy will be followed beyond that decision. Such binding of a decision for each event at each state amounts to a policy, and it turns out that that policy is strictly better than the reference policy unless the reference policy is optimal. While our goal here is not maximization of an objective function, we will use policy improvement to seek a spectrum sharing policy that is strictly better than the lock-out policy in terms of revenue. Existence of such a policy would generally depend on the pair (r_2, λ_2); and whenever it exists we declare (r_2, λ_2) as a point within the profitability region.

As alluded above, carrying out the policy improvement procedure entails knowing the performance of the lock-out policy, but more detailed knowledge beyond R_{LO} is needed: Let $a(\boldsymbol{x})$ be the number of locations that are eligible for granting new requests at state $\boldsymbol{x}$. In other words, $a(\boldsymbol{x})$ is the cardinality of the set $\{i : \boldsymbol{x}+\boldsymbol{e}_i \in S\}$. Let us define the row vector $\boldsymbol{g} = \{g(\boldsymbol{x}) : \boldsymbol{x} \in S\}$ by setting $g(\boldsymbol{x}) = r_1 a(\boldsymbol{x})$. That is, each entry $g(\boldsymbol{x})$ of $\boldsymbol{g}$ is the *instantaneous* rate of revenue generation given that the network is in state $\boldsymbol{x}$. Finally let $\boldsymbol{h} = \{h(\boldsymbol{x}) : \boldsymbol{x} \in S\}$ be a (row) vector that satisfies

$$Q\boldsymbol{h}^T + \boldsymbol{g}^T = \boldsymbol{0}, \quad (3)$$

where Q is the generator matrix of network occupancy under the lock-out policy as defined by (1).

Since row sums of Q are zero Q is not of full rank. Therefore, equality (3) has infinitely many solutions in $\boldsymbol{h}$. The rank of Q is indeed $|S|-1$; and while (3) has many solutions in $\boldsymbol{h}$ the differences $h(\boldsymbol{x}) - h(\boldsymbol{x}') : \boldsymbol{x}, \boldsymbol{x}' \in S$ are unique. Hence (3) has a unique solution in $\boldsymbol{h}$ if an arbitrary entry of $\boldsymbol{h}$, say $h(\boldsymbol{0})$ is set to 0. In turn $\boldsymbol{h}$ can be obtained by inverting a $(|S|-1)\times(|S|-1)$ submatrix of Q.

The vector $\boldsymbol{h}$ has an important interpretation that provides key insight in determining profitability of secondary requests [24, Chapter 8]: The difference $h(\boldsymbol{x}) - h(\boldsymbol{x}')$ is the difference in the total revenue (over an infinite time horizon) if the network is started from state $\boldsymbol{x}$ instead of state $\boldsymbol{x}'$, and from then on it is operated under the lock-out policy. Therefore,

$$h(\boldsymbol{x}) - h(\boldsymbol{x} + \boldsymbol{e}_i)$$

is the opportunity cost of an additional request admitted at location i when the network state is $\boldsymbol{x}$. This cost includes the potential of rejecting primary requests due to the presence of the admitted request, as well as all cascading effects of such rejection that propagate in the network through neighborhood relations.

The policy improvement procedure asserts that admitting a secondary request at location i when the network state is $\boldsymbol{x}$ would improve the revenue if the immediate revenue r_2 of that decision exceeds its opportunity cost. Hence for a given secondary price-demand pair (r_2, λ_2), with $\lambda_2 > 0$ so as to avoid trivialities, there exists a policy that is strictly more profitable than the lock-out policy if and only if

$$r_2 > h(\boldsymbol{x}) - h(\boldsymbol{x} + \boldsymbol{e}_i)$$

for at least one location i and one state $\boldsymbol{x}$. It should be noted that this latter condition involves r_2 (and also r_1, λ_1, G since these determine $\boldsymbol{h}$) but not λ_2.

We summarize this observation by the following theorem:

Theorem VI.1 *Let*

$$r^* = \min_i \min_{\boldsymbol{x} \in S} \{h(\boldsymbol{x}) - h(\boldsymbol{x} + \boldsymbol{e}_i)\}.$$

a) If $r_2 > r^*$ then there exists a spectrum sharing policy that is profitable for any $\lambda_2 > 0$.
b) If $r_2 \leq r^*$ then no spectrum sharing policy yields positive profit, for any λ_2.

The theorem asserts that profitability of spectrum sharing is insensitive to the secondary demand λ_2 and is solely determined by the secondary price r_2. If a secondary price $r_2 > r^*$ generates positive secondary demand then profitability of spectrum sharing is guaranteed by a policy that grants a secondary request at each location i if at the time of the request the system state $\boldsymbol{x}$ is such that $r_2 > h(\boldsymbol{x}) - h(\boldsymbol{x} + \boldsymbol{e}_i)$.

There are two practical impediments to realizing the full profitability region of spectrum sharing. The first issue is computational and concerns calculation of $\boldsymbol{h}$. As alluded above calculation of $\boldsymbol{h}$ entails inverting a square matrix of dimension $|S| - 1$. Recall that $|S|$ is the number of independent sets of the graph G. While the matrix inversion has complexity that is polynomial in $|S|$, for reasonable topologies of G the quantity $|S|$ itself is exponential in the number of locations in G. For example if G is a rectangular grid with n nodes then $|S| = \Omega(1.5^n)$ [21]. More specifically $|S| = 201030$ in the 32 node topology of Fig. 2. Hence calculating $\boldsymbol{h}$ appears intractable even for topologies with moderate size.

The second issue is operational and concerns dynamic requirements of a profitable policy. As described above, in order to admit or reject a secondary request, such a policy may require the entire instantaneous state $\boldsymbol{x}$ of the network. Since each location in a network has its own requests and terminations, the entire network state changes at a rate that is proportional to the number of locations. Hence for large networks availability of instantaneous network state to any entity is unrealistic. This limitation motivates profitability analysis of practical policies that rely on local information to make decisions at each location. The next section studies one such policy, complete sharing, which is appealing for pricing of spectrum access in the private commons regime.

VII. Profitability region of complete sharing

A *complete sharing* policy is specified by a subset L_+ of locations in G. Within L_+ the policy exerts no admission control on secondary requests; hence secondary requests are subject to the same admission criteria as primary requests. The policy flatly rejects all secondary requests in the remaining locations.

For clarity of exposition, in this paper we analyze the complete sharing policy when L_+ is the entire set of locations in G.

We first give an expression for the revenue rate of complete sharing, and then compare that to the reference value R_{LO}. Since secondary requests behave exactly the same way as primary requests once they are granted, the network occupancy under complete sharing (in the entire network) behaves as it does under the lock-out policy, but with request rate $\lambda_1 + \lambda_2$ at each location. Hence in equilibrium the network occupancy has distribution $\pi_{\lambda_1+\lambda_2}$. The analysis of Section V then applies to identify the revenue rate of the complete sharing policy as

$$R_{CS}(r_2, \lambda_2) = \frac{r_1\lambda_1 + r_2\lambda_2}{\lambda_1 + \lambda_2} E_{\lambda_1+\lambda_2}[T]. \qquad (4)$$

Let us define the function $r_{CS} : \mathbb{R}_+ \mapsto \mathbb{R}$ by setting

$$r_{CS}(\lambda_2) = r_1 \left(\frac{E_{\lambda_1}[T]}{E_{\lambda_1+\lambda_2}[T]} - \frac{\lambda_1}{\lambda_2}(1 - \frac{E_{\lambda_1}[T]}{E_{\lambda_1+\lambda_2}[T]}) \right). \qquad (5)$$

It can be verified that

$$R_{CS}(r_{CS}(\lambda_2), \lambda_2) = R_{LO}. \qquad (6)$$

That is, $r_{CS}(\lambda_2)$ is the value of secondary price r_2 that would render secondary demand λ_2 neutral in terms of profit.

Consulting the expression (4) for a fixed value of λ_2 we see that the expectation $E_{\lambda_1+\lambda_2}[T]$ does not depend on r_2 whereas the ratio $(r_1\lambda_1 + r_2\lambda_2)/(\lambda_1 + \lambda_2)$ is increasing in r_2. Hence $R_{CS}(r_2, \lambda_2) > R_{LO}$, and therefore complete sharing is profitable, if and only if $r_2 > r_{CS}(\lambda_2)$. In turn $r_{CS}(\lambda_2)$ delineates the profitability region of complete sharing.

In contrast to the full profitability region identified by Theorem VI.1, profitability region of complete sharing is generally not bounded by a straight vertical line. Yet an analogous, although looser, insensitivity result can be established:

Theorem VII.1 *Let*

$$r^*_{CS} = \max_{\lambda_2 \geq 0} r_{CS}(\lambda_2).$$

a) If $r_2 > r^*_{CS}$ then complete sharing is profitable for any $\lambda_2 > 0$.
b) If $r_2 < r^*_{CS}$ then there exists a value λ_2 such that complete sharing yields lower revenue than the lock-out policy for the secondary price-demand pair (r_2, λ_2).

Proof: Part a) follows since if $r_2 > r^*_{CS}$ then $r_2 > r_{CS}(\lambda_2)$ for any $\lambda_2 \geq 0$. In that case complete sharing generates strictly higher revenue than the lock-out policy if the secondary demand λ_2 is strictly positive so that secondary requests actually exist. If $r_2 < r^*_{CS}$ then due to the definition of r^*_{CS} there is a $\lambda_2 > 0$ such that $r_2 < r_{CS}(\lambda_2)$. This establishes part b). ■

Fig. 3.a illustrates the profitability curve $r_{CS}(\lambda_2)$ of complete sharing for the 32-node topology of Fig. 2 with the primary demand $\lambda_1 = 0.1$ and primary price $r_1 = 1$. The maximum value of this curve is 0.3135 and it occurs at $\lambda_2 = 0$. Hence any secondary price larger than 0.3135 is guaranteed to be profitable if it generates demand. This profitability curve also has a minimum value 0.1769, which occurs in the limit as $\lambda_2 \to \infty$. This implies that any secondary price less than 0.1769 is guaranteed to incur revenue loss, provided that it generates demand. Profitability of price values between 0.1769 and 0.3135 depend on the underlying demand curve.

We also note that since it is possible to incur positive profit beyond the price value 0.1769, the critical price r^* that determines the full profitability region for this example should satisfy $r^* < 0.1769$.

Finally in Fig. 3.b the critical price r^*_{CS} is plotted against the primary demand λ_1, with fixed primary price $r_1 = 1$. The plot shows that the critical price is always less than the primary price, though it is increasing in the primary demand. This behavior can be interpreted as follows: As the primary demand increases, more primary requests are rejected due to granted secondary requests. The increase in the secondary price reflects compensation for the associated opportunity cost.

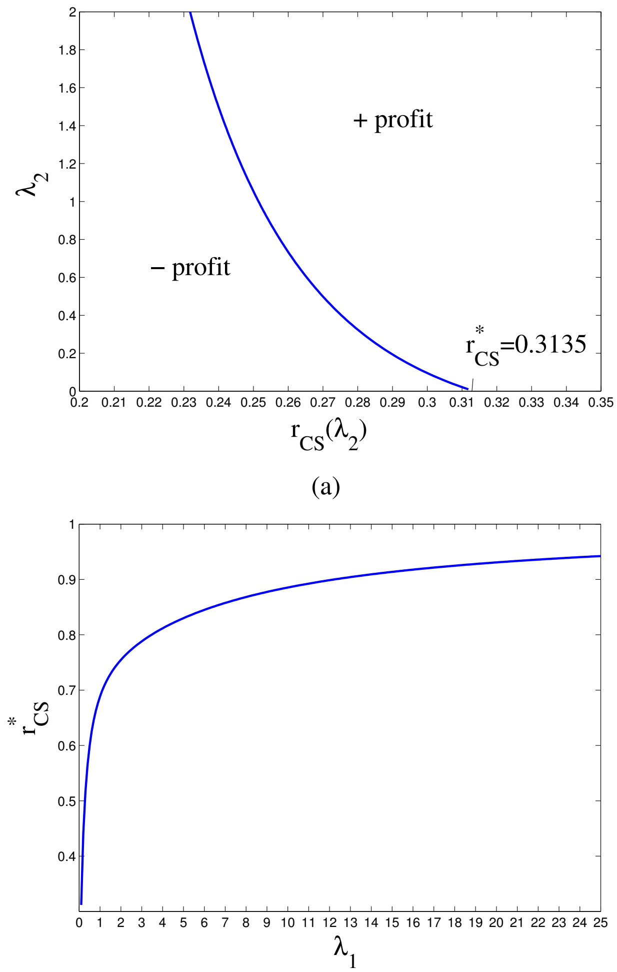

(a)

(b)

Fig. 3. Profitability region and critical prices for complete sharing in the topology of Fig. 2. a) The profitability curve $r_{CS}(\lambda_2)$ for primary price-demand pair $r_1 = 1$, $\lambda_1 = 0.1$. b) For fixed primary price $r_1 = 1$ the critical price r^*_{CS} is increasing in the primary demand λ_1.

VIII. Discussion

a) Larger prices: It is possible to characterize other price thresholds that provide the same profitability guarantee as r^*_{CS}. An example for such threshold is

$$\max_i \max_{\boldsymbol{x} \in S} \{h(\boldsymbol{x}) - h(\boldsymbol{x} + \boldsymbol{e}_i)\} \tag{7}$$

where $h(\boldsymbol{x})$ is defined in Section VI. Note that if r_2 is above that value then the policy improvement procedure of Section VI gives completes sharing as the profitable policy. However, this threshold is overly conservative and in fact the value of (7) can be much larger than the primary price r_1. Theorem VII.1.b asserts that r^*_{CS} is the smallest of such thresholds. This property is important since strict profitability relies on existence of demand, and if r^*_{CS} does not generate demand then neither would a higher price.

b) Protecting the primary demand: In some situations, such as providing secondary service opportunistically in a cellular network, it may be desirable to enhance the revenue while maintaining a certain quality of service level for the primary users. For example, the licensee may aim to grant at least 95% of voice call requests by primary users. Depending on the underlying demand curve the generated secondary demand λ_2 may be too large to maintain such a constraint. In that case the licensee may choose to thin the secondary traffic by declining a fraction of secondary requests. If thinning is applied uniformly over the locations then the overall effect is to reduce λ_2 to an appropriate value so that primary users are protected. Since the secondary price is not altered, positive profit remains guaranteed as long as this reduced demand value is positive. If the quality constraint is satisfied under lock-out policy then there is a positive λ_2 such that it is still satisfied after inclusion of secondary demand.

c) Iterative spectrum offerings: Insensitivity of the profit can be leveraged to increase the revenue further, without the knowledge of the demand curve. Here we outline an iterative strategy:

Let us consider a secondary-spectrum offering at price $r_2 = (1+\varepsilon)r^*_{CS}$ for some $\varepsilon > 0$ that generates demand λ_2. After the offering, the total (primary plus secondary) demand on the spectrum will be $\lambda_1 + \lambda_2$ and the average revenue per request will be

$$\frac{r_1\lambda_1 + r_2\lambda_2}{\lambda_1 + \lambda_2}.$$

Although λ_2 is unknown prior to the offering it can be measured once spectrum sharing takes place; hence both values are known to the licensee after the offering. These values can now be interpreted as primary load and primary revenue respectively, leading to a new critical secondary price $r^{*(2)}_{CS}$. In turn a second round of spectrum offering can be made at price $r_2^{(2)} = (1+\varepsilon)r^{*(2)}_{CS}$. Since part of the potential demand is collected at the first offering, the second offering may be subject to a different demand curve. Yet this new offering will increase the revenue if it generates demand, and will have no adverse effect otherwise.

Following the second offering the spectrum will be carrying demand λ_1 at price r_1, demand λ_2 at price r_2, and some demand $\lambda_2^{(2)}$ at price $r_2^{(2)}$. New values of total demand and average per-request revenue can be easily calculated and adopted in the next iteration of the same procedure. At each iteration revenue is enhanced further if the offering raises additional demand. A block diagram of this approach is given in Fig. 4.

Fig. 5 gives the revenue for several rounds of spectrum offerings on the 32-node graph in Fig. 2. Here we have

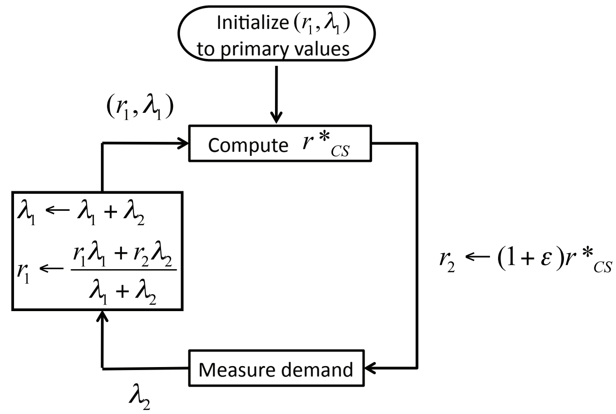

Fig. 4. A block diagram for the iterative procedure to enhance the profit. The goal of the procedure is to raise further demand at each iteration while maintaining strict guarantees on profitability.

made the following choices in modeling the demand curve (it should be noted that the model is used only to determine the demand for an *already decided-upon* price): It is assumed at each round k the secondary demand generated by price r is given by

$$\lambda_2^{(k)}(r) = \int_r^\infty f_k(x)dx$$

for a non-negative function $f_k(x)$. For the initial round $f_1(x)$ can be interpreted as the density of users for whom spectrum access has value x; hence these users will take price r if $r \le x$. The above integral for $k=1$ indicates that the demand generated by price r consists of such users. Assuming that all users with valuation at least r_2 are cleared in the first round, we take

$$f_2(x) = \begin{cases} f_1(x) & \text{if } x < r_2 \\ 0 & \text{else} \end{cases}$$

for the second round. More generally, if we represent the secondary price at the ith offering by $r_2^{(i)}$,

$$f_k(x) = \begin{cases} f_1(x) & \text{if } x < \min\{r_2^{(1)}, r_2^{(2)}, \cdots r_2^{(k-1)}\} \\ 0 & \text{else} \end{cases}$$

for the kth round. To illustrate also the robustness of the outlined procedure against variations in the demand, Fig. 5 gives revenues under two different kernels for the demand curve. The first choice is $f_1(x) = 1$ if $0 \le x \le 1$ and $f_1(x) = 0$ otherwise, thereby reflecting a uniform distribution of spectrum valuation among potential users. The second choice is $f_1(x) = e^{-x}$, reflecting an exponential distribution of spectrum valuations. We point out that the average valuation of spectrum is respectively 0.5 and 1 for these two kernels; so the latter case should be expected to yield higher revenue.

In obtaining the numerical values the primary price-demand values are taken as $r_1 = 1$, $\lambda_1 = 0.1$;

	Round k			
	1	2	3	4
price $r_2^{(k)}$	0.3762	0.3612	0.3610	0.3610
demand $\lambda_2^{(k)}$	0.6238	0.0150	0.0002	0
overall revenue	2.6819	2.6891	2.6892	2.6892

(a)

	Round k			
	1	2	3	4
price $r_2^{(k)}$	0.3762	0.3614	0.3613	0.3613
demand $\lambda_2^{(k)}$	0.6864	0.0102	0.0001	0
overall revenue	2.7186	2.7232	2.7233	2.7233

(b)

Fig. 5. The revenue values obtained by the iterative procedure of Fig. 4 applied on the topology of Fig. 2. The primary price and demand are respectively $r_1 = 1$ and $\lambda_1 = 0.1$. The resulting reference value that determines profitability is $R_{LO} = 2.1227$. A different demand curve is assumed at each iteration as explained in the text. Considered spectrum-valuation densities are: a) $f_1(x) = 1$ for $0 \le x \le 1$. b) $f_1(x) = e^{-x}$.

consequently the reference revenue for profitability is $R_{LO} = 2.1227$. We have also chosen, somewhat arbitrarily, $\varepsilon = 0.2$. It should perhaps be noted that eventual revenue may be smaller or larger depending on the value of ε and also the underlying demand curve, though the monotonic nature of generated revenues is invariant.

IX. Conclusion

This paper investigates economic feasibility of secondary spectrum provision in spatio-temporal settings. Its main contribution is existence and characterization of pricing guidelines that guarantee profitability irrespective of the underlying relationship between price and demand. More specifically we identify critical prices above which positive profit is achieved as long as there is demand. This conclusion may be reassuring for licensees that face challenges due to volatility of market demand. We have also devised a mechanism for spectrum offering that leverage this insensitivity to improve profit.

As opposed to profit-maximization, profitability of dynamic resource sharing has not been studied comprehensively in literature. There seem to be two complementary directions that warrant further study in the context of dynamic spectrum access: The first one concerns implications of the insensitivity property, particularly financial instruments that leverage it. The second direction is more technical and aims to exploit structural properties of the profitability region to mitigate the computational effort required for determining critical prices.

Appendix

Obtaining the revenue rates of both the lock-out and the complete sharing policies, and in turn the curve $r_{CS}(\cdot)$, entails computation of expectations $E_\lambda[T]$ for some $\lambda > 0$. Since $T = \sum_i x_i$, this expectation is

$$E_\lambda[T] = \sum_{\boldsymbol{x}\in S} (\sum_i x_i)\pi_\lambda(\boldsymbol{x}).$$

Hence in principle $E_\lambda[T]$ can be computed by first determining the distribution π_λ and then by computing the above summation. Although it is rather standard, this approach has computational and storage requirements that grow exponentially with the network size: Firstly, computation of the normalizing constant C in the expression (2) for π_λ is NP-hard [22]. Secondly, as alluded earlier, the number of states $|S|$ of $\boldsymbol{x}$ is typically exponential in the number of locations n [21]; hence, even if C is computed, storing π_λ requires memory whose size is exponential in n.

In this section we describe an alternative approach that provides significant complexity advantages by exploiting the structure of π_λ. Towards that end we work with the moment generating function $M(s)$ of T: For each number s, the value of this function is defined as

$$M(s) \ = \ E_\lambda[e^{sT}] \ = \ \frac{\sum_{\boldsymbol{x}\in S}\lambda^T e^{sT}}{\sum_{\boldsymbol{x}\in S}\lambda^T}.$$

Let us define

$$Z(s) = \sum_{\boldsymbol{x}\in S} e^{sT} \tag{8}$$

so that

$$M(s) = \frac{Z(s+\ln\lambda)}{Z(\ln\lambda)}$$

and

$$E_\lambda[T] \ = \ \frac{dM(s)}{ds}|_{s=0} \ = \ \frac{Z'(\ln\lambda)}{Z(\ln\lambda)}. \tag{9}$$

Here Z' represents the derivative of the function Z.

For a given graph G define the vector $\boldsymbol{m} = (m_0, m_1, m_2, \cdots, m_n)$ by setting

$$m_k = \text{number of independent sets of size } k \text{ in } G.$$

The total number of independent sets is then $|S| = \sum_{k=0}^n m_k$. Note that m_k is the number of states $\boldsymbol{x}$ for which $T = \sum_i x_i = k$. Each such state contributes e^{sk} to the summation in (8) so

$$Z(s) = \sum_{k=0}^n m_k e^{sk}.$$

In turn by relation (9)

$$E_\lambda[T] = \frac{\sum_{k=0}^n k m_k \lambda^k}{\sum_{k=0}^n m_k \lambda^k}.$$

Hence $E_\lambda[T]$ can be computed with $O(n)$ arithmetic operations once the vector $\boldsymbol{m}$ is obtained. Note that the size of $\boldsymbol{m}$ is at most $(n+1)$ because an independent set is a subset of the locations and largest independent cannot have size more than n. So, in contrast to π_λ, the memory requirement of storing $\boldsymbol{m}$ scales linearly with the number of locations n. An additional advantage of this approach is that the same vector $\boldsymbol{m}$ can be used to compute $E_\lambda[T]$ for any value of λ; whereas computing $E_\lambda[T]$ via π_λ requires a separate and highly complex calculation for the constant C for each value of λ.

To best of the authors' knowledge counting the number of independent sets *of a given size* has received scant attention in the literature, although there are results that suggest that it entails exponential complexity in the number of nodes [23]. Here we provide a recursive procedure to compute $m_2, m_3, \cdots, m_n$ (note that $m_0 = 1$ and $m_1 = n$). The worst case complexity of this procedure over all graphs G is exponential in n, yet each recursion step is rather simple.

Proposition A.1 *Let $m_k(G)$ be the number of independent sets of size k in graph G. For a given node i of G, let $G-i$ be the subgraph obtained by deleting node i (and all incident edges), and let $G-N(i)$ be the subgraph obtained by deleting i and all neighbors of i (and all incident edges). Then*

$$m_k(G) = m_k(G-i) + m_{k-1}(G-N(i)), \quad k = 1, 2, \cdots \tag{10}$$

with the initial values $m_0(G) = 1$ and $m_k(\emptyset) = 0$ for $k = 1, 2, \cdots$.

Proof: Let us fix an arbitrary node i in G. If an independent set $\boldsymbol{x}$ of G contains i then it cannot contain any neighbors if i. Therefore, if $\boldsymbol{x}$ has k elements (including i) then the remaining elements of $\boldsymbol{x}$ after excluding i must constitute an independent set of size $k-1$ in the subgraph $G-N(i)$. Since $m_{k-1}(G-N(i))$ is the number of such independent sets, it is also the number of independent set with size k in G containing i. On the other hand, if $\boldsymbol{x}$ does not contain i, then it must be an independent set of size k in $G-i$. The number of such independent sets is $m_k(G-i)$. The recursion (10) reflects that $m_k(G)$ is the sum of the number independent sets that include i and the number of independent sets that do not include i. The boundary condition $m_0(G) = 1$ indicates that the only independent set with 0 nodes is the empty set (equivalently $\boldsymbol{x} = \boldsymbol{0}$), and $m_k(\emptyset) = 0$ for $k = 1, 2, \cdots$ indicates that a degenerate graph with no nodes does not possess independent sets other than the empty set. ■

References

[1] "Promoting efficient use of spectrum through elimination of barriers to the development of secondary markets," Second Report and Order, Order on Reconsideration, and Second Further Notice of Proposed Rulemaking (FCC 04-167), September 2004.

[2] J.W. Mayo and S. Wallsten, "Enabling efficient wireless communications: The role of secondary spectrum markets," *Information Economics and Policy*, vol. 22, pp. 61–72. 2010.

[3] O. Ileri, D. Samardzija, and N. Mandayam, "Demand responsive pricing and competitive spectrum allocation via a spectrum server," *IEEE DySPAN 2005*.

[4] J. Huang, R. Berry, and M. Honig, "Auction-based spectrum sharing," *ACM/Springer Journal of Mobile Networks and Applications*, vol. 11, pp. 405–418, June 2006.

[5] J. Bae, E. Beigman, R. Berry, M. Honig, and R. Vohra, "Sequential bandwidth and power auctions for distributed spectrum sharing," *IEEE Journal on Selected Areas in Communication*, vol. 26, pp. 1193–1203, September 2008.

[6] Y. Xing, R. Chandramouli, and C. Corderio, "Price dynamics in competitive agile spectrum access markets," *IEEE Journal on Selected Areas in Communications*, vol. 25, pp. 613–621, April 2007.

[7] X. Zhou and H. Zheng, "Trust: A general framework for truthful double spectrum auctions," *IEEE INFOCOM 2009*.

[8] C.E. Caicedo and M. Weiss, "The viability of spectrum trading markets," *IEEE DySPAN 2010*.

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

[10] H. Mutlu, M. Alanyali, and D. Starobinski, "On-line pricing of secondary spectrum access with unknown demand function and call length distribution," *IEEE INFOCOM 2010*.

[11] O. Besbes and A. Zeevi, "Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms," *Operations Research*, 2009, vol. 57, pp. 1438-1450, November 2009.

[12] D. Bertsimas and G. Perakis, "Dynamic pricing: A learning approach," *Applied Optimization*, vol. 101, pp. 45–79, June 2006.

[13] A. Tonmukayakul and M. Weiss, "Secondary use of radio spectrum: A feasibility analysis," in *The 32nd Research Conference on Communication, Information and Internet Policy*, (George Mason University, Washington, DC), August 2004.

[14] D. Hatfield and P. Weiser, "Property rights in spectrum: taking the next step," *IEEE DySPAN 2005*

[15] J. Chapin and W. Lehr, "The path to market success for dynamic spectrum access technology," *IEEE Communications Magazine*, pp. 96–103, June 2007.

[16] G. Faulhaber and D. Farber, "Spectrum Management: Property Rights, Markets, and Commons," *Telecommunications Policy Research Conference (TPRC) 2003*.

[17] J. Peha, "Approaches to Spectrum Sharing," *IEEE Communications Magazine*, pp. 10–11, February 2005.

[18] W. Lehr and J. Crowcroft, "Managing Shared Access to a Spectrum Commons," *IEEE DySPAN 2005*.

[19] M. Cave, C. Doyle, and W. Webb, Essentials of modern spectrum management. Cambridge, Cambridge University Press, 2007.

[20] J. Peha, "Sharing Spectrum Through Spectrum Policy Reform and Cognitive Radio" *Proceedings of the IEEE*, vol. 97, pp. 708–719, April 2009.

[21] N.C. Calkin and H.S. Wilf, "The number of independent sets in a grid graph," *SIAM Journal on Discrete Mathematics*, vol. 11, no. 1, pp. 54–60 , 1998.

[22] M. Luby and E. Vigoda, "Approximately counting up to four," *Proc. 29th ACM Symp. on Theory of Computing*, pp. 682–687, 1997.

[23] C. Hoffmann, "Exponential time complexity of weighted counting of independent sets," arXiv:1007.1146v2 [cs.CC].

[24] R.A. Howard, Dynamic Programming and Markov Processes, MIT Press, 1960.