

Economic Viability of Private Commons: Framework and Guidelines for Profitability

Abstract

Ongoing regulatory reforms have led to several novel spectrum sharing models under the general umbrella of dynamic spectrum sharing. The private commons model introduced by FCC in 2004 allows spectrum licensees to provide secondary access to spectrum on an opportunistic basis while retaining ownership. Since wireless communication systems are typically overprovisioned in order to deliver service-level guarantees to (primary) users under short-term load variations, this model bears significant potential by facilitating utilization of temporal surplus of capacity through serving secondary users at possibly different service level. A potential barrier to adoption of the private commons model is the uncertainty about secondary price-demand relationship which is difficult to predict in an emerging market: A selected price for secondary access may be profitable for some values of secondary demand but not for others, leading to a profound uncertainty about ultimate benefit of spectrum sharing. This paper aims to eliminate such uncertainty by devising concrete guidelines and methods for profitability. The paper establishes that the price of secondary spectrum access can be chosen to guarantee profitability for any value of secondary demand: It is shown that for both the coordinated and uncoordinated commons regimes a profitable price should exceed a threshold value, which can be calculated. Hence profitability of private commons is insensitive to the relationship between price and demand. The paper also leverages the insensitivity property and provides a technique to further enhance revenue via iterative spectrum offerings.

1. Introduction

Growing demand for wireless spectrum is driving global regulatory reforms that can be collected under the general umbrella of dynamic spectrum access. These reforms consist of several complementary directions towards the primary goal of sustained efficiency of spectrum usage through novel and flexible models. An important direction is secondary spectrum markets which

¹This work was funded in part by the National Science Foundation through grant CCF-0964652.

²This paper was presented in part at IEEE DySPAN 2011.

hinge on extended property rights for license holders and facilitate high spectrum utilization by allowing spectrum ownership to float in response to varying demand and supply conditions.

Spectrum markets may operate at various time-scales. In the slow time-scale, transactions are in the form of long-term spectrum leases for exclusive use. An example of such markets is the spectrum exchange model presently offered by Spectrum Bridge Inc. In the fast time-scale, each transaction spans a shorter duration. In its Second Report and Order on Reconsideration and Second Further Notice of Rulemaking [1] the FCC introduced the *Private Commons* model as a framework for such transactions. In broad strokes, under the private commons model, ownership of spectrum remains with the licensee but the licensee may provide spectrum access to secondary users for a fee. Access may be *coordinated* by the licensee through signals that determine when or how it is allowed, or it may be *uncoordinated* thereby generating perception of commons for secondary users that is akin to the ISM bands. In effect, the model paves the way for spectrum spot-markets for temporal surplus of capacity from any primary usage.

Private commons have significant potential to increase spectrum utilization since communication systems are typically over-engineered in order to provide service guarantees in the face of unpredictable short-term variations in their load. For example, fundamental models of telephony [2] prescribe that to keep the denial of service rate below 5% the average load of a 16 channel spectrum band should be less than 11.50. This implies that about 12 of the 16 channels are busy on average; the remaining channels are typically used only during peak fluctuations of load to maintain the required service level. This observation is supported by a recent measurement-based study over 19,624 GSM base stations deployed in Germany [3], which reports that the majority of base station in crowded areas such as city centers remain underloaded at all times. Hence, in principle, a service provider can enhance its revenue by catering to secondary users at possibly lower service level on aggregate, while at the same time increasing spectrum utilization without compromising its primary service.

Realizing this potential requires that providing spectrum access under the private commons model is ultimately profitable for licensees. A transaction is profitable if its revenue exceeds its cost. In the present context, the revenue of a secondary spectrum access is the associated fee collected by the licensee. The cost of spectrum access seems harder to quantify and requires some analysis: Providing temporal access to spectrum incurs an *opportunity cost* for the licensee. This cost reflects the lost revenue due to the primary demand that could be serviced in the absence of the original access. Clearly this depends on the primary traffic load and it is arguable that the two quantities have a positive correlation. The difficulty of identifying the opportunity cost is further compounded by the spatial dimension of spectrum since effects of interference in fact extend beyond the geographical proximity of the location where spectrum is used. The illustration of Figure 1 helps explain the alluded networking effects: A transmission at the center precludes other spectrum access in zone 1 due to the interference it generates. While this externality appears




Figure 1: Graphical illustration of the example used to explain extended effect of interference.

negative, it would in fact help with spectrum access in zone 2 which would not have been possible had a spectrum access occurred at a geographically close portion of zone 1. Hence, the impact of the original transmission appears positive in zone 2. A similar argument can be made for negative externality in zone 3, and so on.

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 that it generates. 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. Yet this relationship can be of critical importance because a given secondary price may be profitable for some secondary demand values but not for others. For example, a low price may be marginally profitable if it generates low demand but at high demand the spectrum may be monopolized by secondary users at a bargain price, thereby leading to a loss. Conversely, one may ask if a low price may be profitable only if a large volume of secondary users are served; hence if it generates high enough demand but not otherwise. Evidently the spectrum sharing policy has bearing on the answer.

A considerably impactful situation arises if for any chosen value of the secondary access price there is at least one secondary demand value for which spectrum sharing leads to a loss by the licensee. As the underlying demand curve is not explicitly available, the licensee then cannot be sure whether it will have to operate at an unfavorable price-demand regime. This uncertainty may constitute a barrier for the private commons model.

In this paper we show that the price of secondary access can be chosen to guarantee profitability regardless of the demand that it generates. Such a price should exceed a threshold value, which can be calculated. The value of the threshold depends on whether the private commons is coordinated or not. Figure 2 illustrates the pricing guideline under the uncoordinated model where no admission criteria are imposed on secondary access requests except for avoiding in-

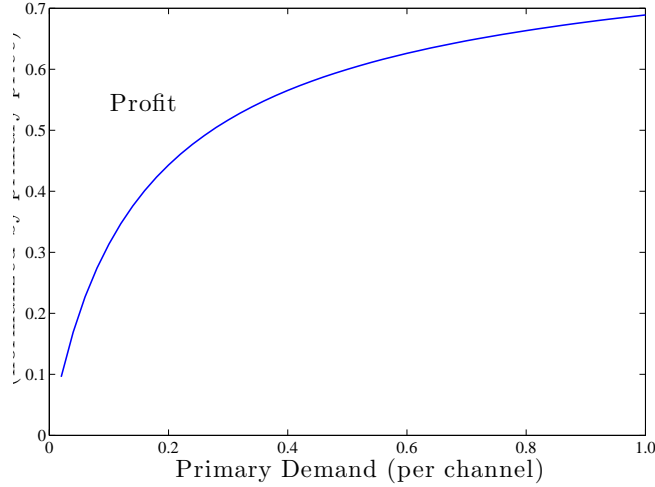


Figure 2: Guideline for profitable pricing strategy for uncoordinated secondary access. The curve indicates critical prices above which secondary access is profitable for the corresponding value of primary demand, regardless of the secondary demand.

interference with existing users. To involve all relevant quantities, the figure shows the threshold after normalizing it with the average revenue obtained from one primary spectrum access. The horizontal axis is primary demand³ per channel -if the spectrum in question is divided into sub-channels. The alluded opportunity cost of secondary access increases with increasing primary demand, and so does the profitability threshold. The threshold price of coordinated commons is substantially smaller but it entails complex management that may be undesirable for practical purposes.

The rest of this paper is structured as follows. Section 2 provides brief pointers to related work. A simple yet insightful spatial model for spectrum access is introduced in Section 3. Section 4 describes profitability of private commons under the considered model and illustrates its interplay with the demand curve. Sections 5 and 6 describe pricing of respectively coordinated and uncoordinated commons in a way to achieve profitability regardless of the demand curve. The conclusions reported in these sections are derived from technical analysis details of which can be found in [15]. In Section 7, we also provide a pricing technique that leverages the insensitivity of profitability to further enhance the revenue. The paper concludes with final remarks in Section 8.

2. Background

The notion of using unlicensed commons for allocating spectrum bands has been embraced as one alternative to the traditional command-and-control model. Another alternative is the property

³Throughout the paper the unit of demand is Erlangs, i.e. average number of placed requests per unit time. Hence the term refers to a scalar number that implicitly accounts for randomness and fluctuations in instantaneous spectrum load.

rights model initially devised by Coase in his seminal paper in 1959 [4]. While there seems to be consensus among the regulatory community on the need to liberate spectrum bands, a debate is taking place on choosing the proper model. For example, in the course of defending the property rights model, Hazlett and Leo argue that uncoordinated interference can lead to conflict between users and only liberal licenses can improve social coordination [5]. In the same time, proponents of spectrum commons establish strong correspondence between the command-and-control model and the property rights model [6]. Their advocacy is influenced by the successful experience with the ISM band and the widespread use of the Wi-Fi technology. In this regard, the private commons model appears to be a suitable model that takes from the two worlds. A taxonomy of the different types of commons is provided by Buddhikot in [7] which promotes private commons as one coordinated type that allows for enforcing access policies.

There is a bulk of work in the literature that address challenges which impede successful implementation of different spectrum sharing models. For example, Lehr et al. detail issues in implementing spectrum commons that range from power control to privacy and security issues [8, 9]. Tonmukayakul and Weiss address the dimensionality nature of radio spectrum (e.g., space, time, frequency) that need to be accounted for in any viable implementation [10] and Peha describes pros and cons of various spectrum sharing mechanisms [11, 12].

However, the authors of the paper are not aware of fundamental analysis on economic viability of private commons as a joint economic and engineering problem. A notable exception perhaps is [13], which studies conditions for market liquidity in exchange-based spectrum trading. A salient feature of the present paper is to provide licensees with tangible guidelines for profitability based on realistic technical models. In this respect, the paper presents a quantitative account for evaluating the opportunity cost for using spectrum.

3. Spatio-temporal Model of Spectrum

Spatial properties of spectrum sharing can be conveniently modeled using a graph representation. An example of such a graph, which has been popular in analysis of cellular networks, is illustrated in Figure 3. In this representation, each circle (vertex) represents a geographical location. An edge is drawn between two circles if the same band of spectrum cannot be utilized simultaneously at the two locations represented by the vertices, due to interference. In that case the two locations are designated as *neighbors*. Thus, if a band is in use at one location it cannot be used at *any* neighboring location.

In a cellular service context, each vertex of the graph may represent a cell with an associated base station. Alternatively, in a general private commons context, the graph may represent a suitable discretization of space. In that case, the number of edges per node in the graph would be increasing with the resolution of the spatial representation. The analytical discussion of the




Figure 3: The hexagonal lattice topology used in modeling spatial aspects of spectrum.

paper does not make any assumptions on the topology of the graph, though all numerical values that are reported in the paper are based on the popular hexagonal lattice in Figure 3.

Spectrum demand studied in this work is composed of two types of requests; *primary* and *secondary*. At each location, requests of either type arrive randomly. We will denote the average number of primary requests that arrive per unit time by λ_1 . Similarly λ_2 denotes the same quantity for secondary requests. If a request at some location is granted then it holds spectrum at that location for unit time on average. During that time spectrum is not available to incoming requests at that location as well as at neighboring locations and such requests are denied.

For the spectrum licensee, a granted request generates revenue r_1 if it is primary, and revenue r_2 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. The licensee does not generate revenue from denied requests.

Primary requests are identified with legacy users of the spectrum provider where it is assumed that such requests are granted whenever possible. That is, a primary request at a certain location is granted if at the time of the request all neighboring locations (locations that share no edge with that location) are idle. 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*. Hence, a secondary request may possibly be denied even if interference requirements do not dictate so.

4. Profitability Definitions

Profitability from a given 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 therefore the




Figure 4: 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.

revenue of what is called a *lock-out policy* that flatly rejects all secondary requests. For a spectrum sharing policy to be profitable, revenue of the policy should exceed the revenue obtained using the lockout policy.

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. In this context, the parameters that determine the rate at which the lock-out policy generates revenue in the long-term are primary demand and primary price.

Revenue rate from a given sharing policy is the aggregate rate of revenue from primary and secondary requests. An increase in spectrum demand due to accommodating secondary requests does not necessarily imply an improvement in revenue. Whether such an improvement arises under a given spectrum sharing policy may depend on the raised secondary demand.

The illustration of Fig. 4 helps clarify the generic relationship between profitability and demand curve. The horizontal and vertical axes of the figure respectively represent the stipulated prices and corresponding demand for 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. 4. 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.

In this context, we define the *full profitability region* as the secondary price-demand pairs (r_2, λ_2) for which there exists some spectrum sharing policy that is profitable. Some reflection verifies that this region is the union of profitability regions of all possible spectrum sharing policies. Interestingly, as will be argued in Section 5, there is a single policy that is profitable for all such price-demand pairs.

5. Profitability of Coordinated Commons

In principle, full profitability region would be identified by listing all possible spectrum sharing policies and by taking the union of their profitability regions. Such a task is evidently intractable but one can instead adopt the so-called *policy improvement* optimization technique as a succinct alternative [16]. 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 at each state of the network, under the assumption that the reference policy will be followed beyond that decision. Such binding of a decision for each state amounts to a policy. 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, policy improvement can be used 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 considered secondary price and demand; and whenever it exists we declare the pair as a point within the profitability region.

Application of policy improvement to the present setting [15] yields that at each network state there is a separate threshold value for each location such that if secondary price exceeds this threshold then it is better to accept the secondary request at that location than rejecting it. Otherwise the request should be rejected. Interestingly, the threshold does not depend on the intensity of secondary requests, i.e. the secondary demand. Since all secondary requests are charged the same, the minimum of these thresholds, call it r^* , determines the a full profitability region: If secondary price r_2 chosen such that $r_2 < r^*$ then under no situation accepting a secondary request is better than rejecting it. Hence the lockout policy is optimal for such prices. But if $r_2 > r^*$ then there is at least one situation for which accepting a secondary request is better; therefore the lockout policy can be improved and by definition the improved policy is profitable.

The full profitability region can be characterized as in Figure 5. This region is delineated by a straight line defined by the horizontal intercept r^* . Hence if $r_2 > r^*$ then profitability is guaranteed for any value of secondary demand λ_2 as long as $\lambda_2 > 0$. Clearly if no secondary

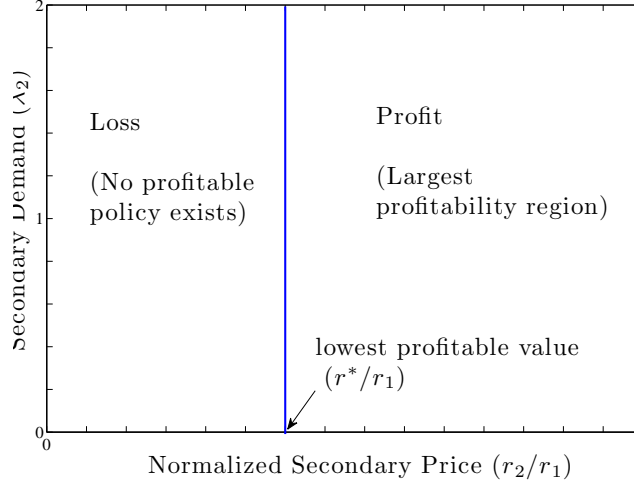


Figure 5: Graphical illustration of full profitability region. If secondary demand is below r^* then there exists no profitable sharing policy. Conversely, there is a spectrum sharing policy that is profitable for any price above r^* so long as the price generates demand.

demand exists then all policies yield the same revenue as the lockout policy, but in that case generating demand by lowering the secondary price only incurs loss.

There are two practical impediments to realizing the full profitability region of spectrum sharing. The first issue is computational: To give a concrete example, computing the secondary price threshold and the policy that achieves the full profitability in the spatial model of Figure 3 requires inverting a matrix of dimension roughly $200K \times 200K$. The second issue is operational and concerns dynamic requirements of the profitable policy. As described above, in order to admit or reject a secondary request, such a policy may require the entire instantaneous state of the network. Since each location in a network has requests and terminations, the entire network state changes fast. Hence availability of the entire instantaneous network state at any instant of time is unrealistic. This limitation is profound and motivates profitability analysis of practical policies that rely on local information to make decisions at each location.

6. Profitability of Uncoordinated Commons

In the uncoordinated private commons regime there are no additional admittance criteria on secondary users except for avoiding interference with existing users at the time of the request. Hence, primary and secondary requests are treated equally. This is the model adopted, for example, in accessing the ISM bands. In this model primary and secondary users differ only due to the revenue they generate, i.e. r_1, r_2 . A fraction $\lambda_1/(\lambda_1 + \lambda_2)$ of the aggregate requests are primary




Figure 6: Profitability region of uncoordinated private commons with primary demand $\lambda_1 = 0.1$ in the hexagonal lattice topology of Figure 3.

and pay r_1 units if accepted; and fraction $\lambda_2/(\lambda_1 + \lambda_2)$ are secondary and pay r_2 upon accessing spectrum.

The revenue of the uncoordinated commons is the same as the lockout policy, but when the demand is $\lambda_1 + \lambda_2$ (rather than λ_1) and the revenue per-accepted request is

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

(rather than r_1). Hence the two policies can be compared relatively easily, and profitability conditions can be recovered.

Profitability region of uncoordinated private commons is delineated by secondary price-demand points that render secondary demand neutral in terms of profit (i.e., revenue rate from uncoordinated access is equal to the rate from the lockout policy). Such points comprise what we refer to as *profitability curve*. Fig. 6 illustrates the profitability region for the hexagonal lattice topology of Fig. 3 with primary demand $\lambda_1 = 0.1$. Here the secondary price is normalized and what is plotted is in fact the ratio of the secondary price r_2 to primary price r_1 . The maximum neutral price is 0.3135 and it occurs at $\lambda_2 \approx 0$. Thus, 31.35% of the primary price is a critical secondary price that is insensitive to secondary demand. Namely, if a price is chosen higher than the critical price, then the uncoordinated commons model is profitable as long as the selected price generates demand.

This profitability curve also has a minimum value 0.1769, which occurs in the limit as λ_2

becomes very large. This implies that any secondary price less than 17.69% of the primary price is guaranteed to incur revenue loss, provided that it generates demand. Profitability of normalized secondary price values between 0.1769 and 0.3135 depends on the underlying demand curve.

6.1. *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. While the uncoordinated model is appealing due to its simplicity it does not support sustaining any service level guarantees for primary users. However a mild overlay policy may be adopted to achieve such guarantees without making any significant sacrifice from simplicity. Namely, the licensee may choose to thin the secondary traffic by declining a fraction of secondary requests upfront, before admitting to the uncoordinated access regime. 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 service level constraint is satisfied by a comfortable margin under the lock-out policy, then there is a positive λ_2 such that it is still satisfied after inclusion of secondary demand; and therefore a window of opportunity to profit from secondary spectrum provision while sustaining service level for primary users.

6.2. *Larger prices*

It is possible to find other price thresholds that provide the same profitability guarantee as those characterized in the paper. Perhaps an immediate example for such threshold would be the primary price r_1 . However, these latter thresholds are overly conservative. The thresholds in Sections 5 and 6 in fact identify the *smallest* profitable prices possible. This property is important since strict profitability relies on existence of demand, and if the threshold price does not generate demand then neither would a higher price. In other words, artificially inflating secondary price may raise no demand at all and it may thereby preclude profitability.

7. **Further Enhancements: Iterative Spectrum Offerings**

The insensitivity of the profit to secondary demand can be leveraged to increase the revenue further, without knowledge of the demand curve. This can be achieved by offering spectrum in rounds according to the following approach:

Consider a spectrum offering for secondary access at a price that exceeds the critical secondary price. Thus, this price is profitable for any secondary demand. To avoid any trivialities it

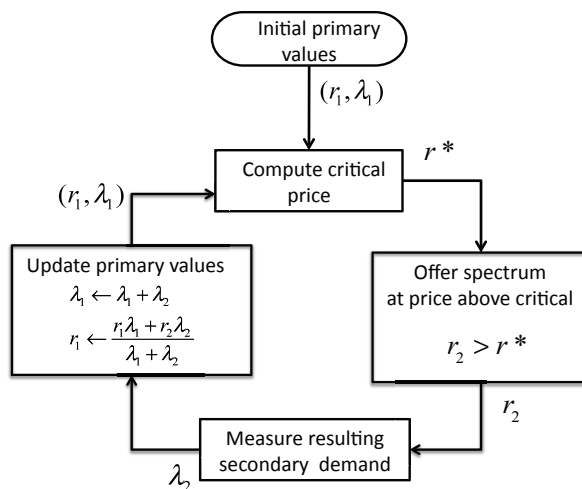


Figure 7: A block diagram for the iterative procedure to enhance the profit. The goal of the procedure is to raise further demand at each iteration while maintaining strict guarantees on profitability.

is implicitly assumed that such a price generates strictly positive demand $\lambda_2 > 0$. Although rate of secondary demand is unknown prior to the offering it can be measured once spectrum sharing takes place; hence both primary and secondary demand rates are known to the licensee *after* the offering. The new request rate becomes $\lambda_1 + \lambda_2$ after this offering and the average revenue per request becomes $\frac{r_1\lambda_1 + r_2\lambda_2}{\lambda_1 + \lambda_2}$.

Now consider a second offering in which these values can be interpreted as primary load and primary revenue, respectively. This interpretation leads to a new secondary price for the new primary parameters; and in turn to a new value for critical secondary price for profitability. This second offering may be subject to a different demand curve since part of the potential demand is already collected at the first offering. Yet if the price is chosen above the critical then the new offering is guaranteed increase the revenue if it generates demand, and it will have no adverse effect on the revenue otherwise.

Following the second offering the spectrum will be carrying (i) primary demand at rate λ_1 and price r_1 , (ii) secondary demand at rate λ_2 and price r_2 after the first offering, and (iii) additional secondary demand at rate $\lambda_2^{(2)}$ and price $r_2^{(2)}$ after the second offering. 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. 7.

Revenue values for several rounds of spectrum offerings are given in Figure 8(a) for the hexagonal lattice topology in Fig. 3. The demand curve for this example is modeled as illustrated in Figure 8(b) where positive demand is generated for prices in the range $(0, 1]$. In the first

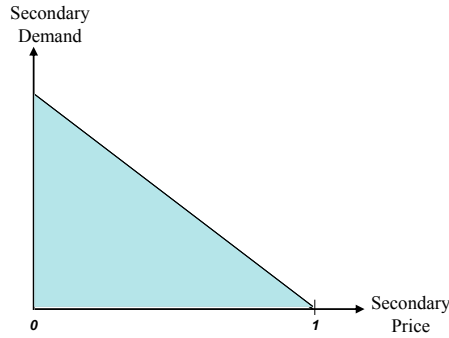
iteration, demand is composed of users whom the value of spectrum access exceeds $r_2^{(1)}$. These users are cleared during the first round as shown in Figure 8(c). Thus, in the second round users for whom spectrum value is between $r_2^{(1)}$ and $r_2^{(2)}$ are going to access spectrum as shown in Figure 8(d). More and more users are going to be cleared at each iteration until the price offered generated no demand; i.e., no users left who value spectrum beyond that price which is found here to be 0.361.

8. Conclusion


The paper studies profitability from dynamic spectrum provision in private commons. The analysis in the paper is aimed to quantify the opportunity cost endured by the licensee from sharing spectrum; i.e., lost revenue due to the primary demand that could be serviced in the absence of the original access. The paper discusses complexities of evaluating this quantity that stem from (i) spatio-temporal effect of interference and (ii) lack of knowledge of price-demand relationships. In this regard, the paper suggests pricing guidelines for balancing the opportunity cost. In particular, the paper shows that, for both the coordinated and the uncoordinated regimes, a price threshold can be specified to guarantee profitability regardless of the demand it generates. It is argued that the threshold price of coordinated commons is substantially smaller but entails complex management. Finally, the paper devises a technique for revenue improvement via iterative offerings.

	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) Secondary price, demand, and revenue at different rounds of spectrum offering for the 32-node graph in Fig. 3. In obtaining the numerical values the primary price-demand values are taken as $r_1 = 1$, $\lambda_1 = 0.1$; consequently the reference revenue for profitability is $R_{LO} = 2.1227$.



(b) Demand curve for secondary spectrum access used in obtaining the results in Figure 8(a).



(c) Users who value spectrum more than $r_2^{(1)}$ are cleared after the first offering of spectrum.

(d) Users who value spectrum between $r_2^{(1)}$ and $r_2^{(2)}$ are cleared after the second offering of spectrum.

Figure 8:

References

- [1] Federal Communications Commission, "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] ITU-D SG 2/16 & ITC. Teletraffic Engineering Handbook. 2001.
 - [3] M. Michalopoulou, J. Riihijärvi and P. Mähönen, "Towards Characterizing Primary Usage in Cellular Networks: A Traffic-based Study," *IEEE Dynamic Spectrum Access Networks DYSPAN 2011*, Aachen Germany, May 2011.
 - [4] R. Coase, "The Federal Communications Commission," *Journal of Political Economy* II (2), 1959.
 - [5] T. W. Hazlett and E. T. Leo "The case for liberal spectrum licenses: A technical and economic perspective," 25 *Berkeley Technology Law Journal* (forthcoming, 2011).
 - [6] L. Lessig, *The Future of Ideas: The Fate of the Commons in a Connected World*. Random House, 2001.
 - [7] M. M. Buddhikot, "Understanding dynamic spectrum access: models, taxonomy and challenges" *IEEE Dynamic Spectrum Access Networks DySPAN 2007*, pp. 649–663, April 2007.
 - [8] W. Lehr and J. Crowcroft, "Managing Shared Access to a Spectrum Commons," *IEEE Dynamic Spectrum Access Networks DySPAN 2005*.
 - [9] J. Chapin and W. Lehr, "The path to market success for dynamic spectrum access technology," *IEEE Communications Magazine*, pp. 96–103, June 2007.
 - [10] 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.
 - [11] J. Peha, "Approaches to Spectrum Sharing," *IEEE Communications Magazine*, pp. 10–11, February 2005.
 - [12] J. Peha, "Sharing Spectrum Through Spectrum Policy Reform and Cognitive Radio" *Proceedings of the IEEE*, vol. 97, pp. 708–719, April 2009.
 - [13] C.E. Caicedo and M.B.H. Weiss, "The viability of spectrum trading markets," *IEEE Communications Magazine*, vol. 49, issue:3, pp. 46–52. 2011.
 - [14] Federal Communications Commission, "In the matter of principles for promoting the efficient use of spectrum by encouraging the development of secondary markets," December 2000.
 - [15] M. Alanyali, A. Al Daoud, and D. Starobinski, "Profitability of Dynamic Spectrum Provision for Secondary Use," *IEEE Dynamic Spectrum Access Networks Dyspan 2011*, May 2011.
 - [16] R.A. Howard, *Dynamic Programming and Markov Processes*, MIT Press, 1960.