Market Structures for Wireless Services with Shared Spectrum

Randall Berry, Michael Honig,Vijay Subramanian
Northwestern University
{rberry, mh, vjsubram}@eecs.northwestern.edu

Thành Nguyen
Purdue University
nguye161@purdue.edu

Rakesh Vohra
University of Pennsylvania
vohra@seas.upenn.edu

Abstract—**There has been much recent interest in allowing commercial providers to utilize federal spectrum on a shared basis. In such scenarios, the use of** *secondary licenses* **has been proposed as a way to provide service providers with greater predictability about the service that can be offered and thus help encourage investment. However, since the spectrum is shared, such a license would differ from traditional exclusive use licenses in that the given spectrum may only be intermittently available. In this work we discuss several issues related to such intermittent spectrum and in particular consider different market structures for allocating such secondary licenses to different service providers. In each case we characterize the complexity of the underlying efficient allocation problem and discuss market mechanisms for implementing such an outcome.**

I. INTRODUCTION

Driven in part by projections of continuing rapid growth in the demand for wireless data services, wireless spectrum policy has been attracting much attention. In particular, many have argued that additional spectrum is need to support this demand. However, nearly all spectrum suitable for wireless data services is currently allocated to either commercial or government uses. One approach is to clear allocated spectrum by moving the current services to different bands, but such band clearing can be prohibitively expensive. An alternative solution in these cases is for wireless service providers to *share* the band with the incumbent users. Recently, there has been a push for such sharing to take place between wireless operators and federal agencies. For example, one of the main recommendations in a 2012 report on wireless spectrum by the President's Council of Advisors on Science and Technology (PCAST) in the United States was to share 1000 MHz of federal spectrum [1]. There have also been Presidential Memorandums in 2010 and 2013 directing that federal spectrum be identified that is suitable for sharing with commercial operators [2], [3] and the Federal Communications Commission has begun considering sharing of some such bands such as the 3.5 GHz band [4].

There has been much work on spectrum sharing in recent years and in particular on sharing by *secondary users* that seek to avoid interfering with primary (incumbent) spectrum users. This includes work on utilizing spectrum sensing by secondary users (e.g. [5]) and work on various market mechanisms that can be used to allocate spectrum to secondary users when the primary user is not active (e.g. [6]). In much of this work it is assumed that secondary use is open, meaning that anyone can use a given band (or participate in a market for that band). A potential issue with such open secondary access is that it may be difficult for a provider to predict the long-term service quality that it can offer in a given band since this will depend in part on which other secondary users seek to use this band in the future. This uncertainty can translate into reluctance of a service provider to invest in offering service in the shared band. One way to better encourage such investment is to enable a provider to obtain an exclusive license to operate as a secondary user in a given band. Such an approach has been proposed by Qualcomm under the name Authorized Shared Access (ASA) [7] for sharing the 3.5GHz band with small cells. The Radio Spectrum Policy Group in the European Union has proposed a similar approach called Licensed Shared Access (LSA) [8]. Such secondary licenses could be granted over a longer time-scale and thus provide secondary users with a more predictable service.

In this paper, we consider a scenario in which such secondary licenses are used to share spectrum with primary users that operate intermittently in time.[1] In such cases, when the primary user is active in a given area we assume that the secondary license holder can not utilize a band; otherwise, the secondary license holder has the exclusive right to access the band. In other words, the secondary license holder has the exclusive right to an intermittently available resource. Given such a resource, we consider the following basic question: *If there were many such bands available, how should the bands be allocated to different service providers?* In particular we focus on a case where different providers seek to acquire bands of intermittent spectrum with different probabilities of being available. Each provider wants one band with a given minimum availability. We consider several different approaches for how such bands may be allocated to users. For each approach we consider the efficient allocation of the bands and also discuss market mechanism that could be used to obtain these.

---

[1]Such licenses can also be used to share spectrum spatially with a primary user; in such cases in an area where the primary is never present, the license would essentially be the same as an exclusive license.

II. Model

We consider a model in which there are K potential channels (i.e., bands of spectrum) available for shared access and a set of I service providers (SPs) that can potentially use these K bands. We assume that time is slotted and in every period, each channel k becomes available with a known probability p_k, independent of the availability of other channels and its own availability in other periods. Such a probability reflects a commitment on the part of the primary user to not utilize the spectrum more than this.[2] If available then all channels are identical from the view of a SP, i.e., all have the same bandwidth and similar propagation characteristics. We assume that a service provider is given a license to a shared channel over a long time-scale, i.e., this time-scale is much longer than the time-scale at which channel availability changes. Assuming that this time-scale is long enough that the law of large numbers applies, then each channel k will be available approximately p_k fraction of the time during the license period.

We further assume that each SP i wants exactly one band of spectrum with a given availability q_i and values such a channel at v_i. For example, q_i may represent the minimum availability for which it is worthwhile for SP i to invest in service and v_i could be the anticipated profit from such investment. Even if a channel has an availability greater that q_i, we assume that its value to the service provider is still v_i, i.e., we consider this as excess supply that the provider is not able to utilize.

In the following we will consider several different markets for such a setting.

A. A simple matching market

The most basic model we consider is one in which at most one channel is allocated to any provider. The socially optimal allocation in this case is the allocation that maximizes the sum of the v_i's over those providers who are allocated a channel. This is simply a maximum weighted matching problem in a bipartite graph. Specifically, as shown in Figure 1, we construct a bipartite graph in which one set of nodes corresponds to the channels and one set corresponds to the SPs. Each SP i is connected to channel k if and only if $q_i \leq p_k$, i.e., the channel satisfies the SP's needed availability. We then put a weight of v_i on each edge from SP i. Any of the well-known algorithms for finding a maximum weighted matching (e.g., the Hungarian algorithm) can then be applied to compute the efficient allocation. Market design for such matching settings have been well-studied and can be applied here. For example, [9] gives a dynamic auction procedure that constructs market clearing prices in which the price for each band increases whenever more than one SP desires it. Also, since determining the efficient allocation is computationally tractable, the well-known Vickrey-Clarke-Groves (VCG) mechanism could also be used.

[2]In some cases predicting such availability for federal spectrum may be difficult, but without some such commitment, secondary users may be unwilling to invest.
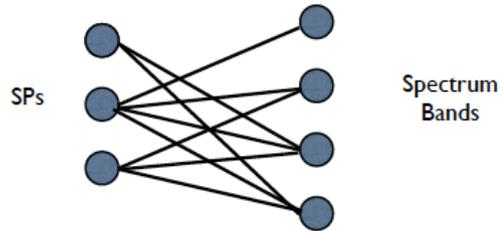

Fig. 1. Bipartite graph between SPs and Spectrum Bands.

B. Channel pooling

Even though SPs only desire a single channel in our model, it may still be beneficial to assign them multiple channels to benefit from statistical multiplexing. For example, consider two channels with availabilities p_k and p_j. If SP i is assigned both of these and is able to dynamically switch to which ever one is available, then it would be able to effectively have one channel with availability $1 - (1 - p_k)(1 - p_j)$. If we allow SPs to pool an arbitrary number of channels, the efficient allocation is then given by

$$\max \sum_i v_i y_i$$

$$\text{s.t.} \sum_k x_{ik} \log(\frac{1}{1 - p_k}) \geq y_i \log(\frac{1}{1 - q_i}) \; \forall i \quad (1)$$

$$\sum_i x_{ik} \leq 1 \; \forall k$$

$$x_{ik} \in \{0, 1\}, y_i \in \{0, 1\}.$$

Here x_{ik} indicates that channel k is assigned to SP i and y_i indicates that SP i is assigned enough channels to meet its needed availability. The first constraint arises since to meet SP i's availability, it must be that

$$1 - \prod_k (1 - p_k)^{x_{ik}} \geq q_i$$

which is equivalent to

$$\prod_k (1 - p_k)^{x_{ik}} \leq (1 - q_i)^{y_i}$$

for $y_i = 1$. The constraint then follows from taking the negative logarithm of both sides of this equation. The second constraint ensures that each channel is assigned to at most one SP. Note that in this problem, we are assuming that the investment costs of a provider do not scale if it uses more than one channel. If this was not the case, then v_i and/or q_i may change with the number of channels allocated.

The problem in (1) is NP hard as it contains the bin covering problem [10] as a special case. Given a collection of items of size at most 1, the goal of the bin covering problem is to partition them into sets so as to maximize the number of sets whose total sum is at least 1. Each such set representing a bin of size one; the bin is covered when its capacity is exceeded. Given an instance of bin covering with a set of items of size s_k each, choose p_k so that $s_k = \log(\frac{1}{1 - p_k})$ for

all k, $v_i = 1$ and $q_i = 1 - \frac{1}{e}$ for all i. The first constraint in (1) then becomes

$$\sum_k x_{ik} s_k \leq y_i.$$

Hence, y_i can be viewed as indicating where the ith bin is covered and x_{ik} indicates if the kth item is placed in that bin. Given that determining the efficient allocation is NP hard, this then implies that designing an efficient market for this setting would also be a hard problem. For example, implementing the VCG mechanism would require being able to solve (1) multiple times.

C. Channel and user pooling

The previous market allows SPs to pool multiple channels, but still requires each channel to be assigned to exactly one SP. Under our assumption that each SP only needs one channel, it follows then if an SP is assigned multiple channels and more than one is available, then the extra channels are "wasted" in the given slot. Next we consider a market design that improves on this by pooling both users and channels. Specifically, we consider a setting where all the channels are pooled and dynamically assigned to the SPs over time based on availability. Each SP is still guaranteed to get one channel with the desired availability, except now this is not a fixed channel but rather a "virtual channel" given to it by an allocation policy. Implementing such a scheme would require some entity to do this allocation in each period. Moreover, to reduce investment costs, it may make sense for such an entity to also invest in common infrastructure that it could then allow each SP to utilize, i.e., the SP could be simply leasing resources on this common infrastructure.

For such a model one then needs to specify a scheduling policy u which determines which SPs are assigned to which channels at each time. Such a rule can be based on the available channels at that time as well as the history of previous assignments and could furthermore be randomized. The main restriction we place on such policies is that they never leave a channel unassigned if there is an agent to which it could be assigned (i.e. an agent without a channel). Any such policy is called *admissible*. For each user i, let $a^u(i)$ denote the availability user i receives under policy u, and for a given subset of users A, let

$$a^u(A) = \sum_{i \in A} a^u(i)$$

be the sum availability of the subset. It is then easy to see that the vector of availabilities satisfies a *strong conservation law* as defined in [11]. Namely, the sum of the availabilities over all of the SPs is the same for all admissible policies, and for any subset A, $a^u(A)$ is maximized by giving A strict priority over all other SPs, i.e,. the users in A are always the first to be allocated a channel. It follows from [11] that the possible performance space is a polymatroid, $\mathcal{P}$, whose vertices correspond to the performance achieved by strict priority policies, i.e., policies in which the SPs are ordered and at each time if L channels are available, they are given to the first L users in this ordered list. For example, if user i is given strict priority in obtaining a channel whenever one is available, it will have availability

$$a_i = 1 - \prod_k (1 - p_k).$$

The efficient allocation in this setting is then given by

$$\begin{aligned}
\max \ & v_i y_i \\
\text{s.t.} \ & a_i > y_i q_i \\
& (a_1, \ldots, a_n) \in \mathcal{P} \\
& y_i \in \{0, 1\}.
\end{aligned} \quad (2)$$

The first constraint here ensures that each user that is allocated a virtual channel receives its desired availability, and the second constraint ensures that these availabilities are in the polymatroid which determines the feasible performance space. This problem is again NP hard as it contains the 0-1 knapsack problem as a special case. Specifically, suppose that there is just one channel and that channel has an availability of $p_1 = 1$. In this case the polymatroid $\mathcal{P}$ becomes a unit simplex, i.e. any non-negative availabilities are achievable if and only if

$$\sum_i a_i \leq 1.$$

If we introduce new variables $z_i = a_i / q_i$, then this can be re-written as

$$\sum_i q_i z_i \leq 1,$$

and if a SP's desired availability is achieved it must be that $z_i = 1$. Thus (2) can be re-written as

$$\begin{aligned}
\max \ & v_i z_i \\
\text{s.t.} \ & \sum_i q_i z_i \leq 1 \\
& z_i \in \{0, 1\}.
\end{aligned} \quad (3)$$

This is exactly the 0-1 knapsack problem.

D. Channel and user pooling with flexible demands

Our final variation is to still allow channel and user pooling as in the previous section but to change our assumption on the user demands. Specifically, we now assume that instead of desiring one channel with a given fixed availability, each SP instead receives a utility that is proportional to the availability of the channel, i.e., if it gets a virtual channel with availability a_i, then its utility is $v_i a_i$. In this case the efficient allocation becomes

$$\begin{aligned}
\max \ & v_i a_i \\
\text{s.t.} \ & (a_1, \ldots, a_n) \in \mathcal{P},
\end{aligned} \quad (4)$$

i.e., this is simply maximizing a linear function over a polymatroid, which can be easily solved via a greedy algorithm. Namely, we assign users priority according to the value v_i, with large values getting higher priority. The availabilities are then given by the vertex of $\mathcal{P}$ that corresponds to that priority order. Furthermore, an efficient ascending auction for this setting can be found in [12].

III. Conclusions

We described several different approaches for allocating intermittent spectrum. In particular we focused on the role of pooling across channels and users as ways to improve the availability of such spectrum. Interestingly, just pooling channels makes the market design problem hard, while pooling over both channels and users makes this problem easier, at least for users with flexible demands. Of course the models we considered here were highly simplified and could be enriched in a number of ways, such as modeling dependencies in the availabilities over time and allowing for more general demands on the part of the providers. There are also many other issues related to sharing government spectrum that we did not touch upon, such as how would service providers value this spectrum and how can federal agencies be incentivized to share their spectrum and commit to a given availability.

References

[1] PCAST, *Realizing the Full Potential of Government-Held Spectum to Spur Economic Growth*. President's Council of Advisors on Science and Technology, July 2012.

[2] "Presidential memorandum: Unleashing the wireless broadband revolution," June 2010. [Online]. Available: http://www.whitehouse.gov/the-press-office/presidential-memorandum-unleashing-wireless-broadband-revolution

[3] "Presidential memorandum – expanding America's leadership in wireless innovation," June 2013. [Online]. Available: http://www.whitehouse.gov/the-press-office/2013/06/14/presidential-memorandum-expanding-americas-leadership-wireless-innovatio

[4] F. C. Commission, "Amendment of the commission's rules with regard to commercial operations in the 3550-3650 mhz band, docket no. 12-354," Notice of Proposed Rulemaking, Dec. 2012.

[5] Q. Zhao and B. Sadler, "A survey of dynamic spectrum access," *IEEE Signal Processing Magazine*, vol. 24, no. 3, pp. 79–89, May 2007.

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

[7] Qualcomm, "LTE advanced: Authorized shared access," 2013. [Online]. Available: http://www.qualcomm.com/research/projects/lte-advanced/asa

[8] Radio Spectrum Policy Group 2011, "RSPG11-392: Report on Collective Use of Spectrum (CUS) and other spectrum sharing approaches," november 2011.

[9] G. Demange, D. Gale, and M. Sotomayor, "Multi-item auctions," *The Journal of Political Economy*, pp. 863–872, 1986.

[10] S. Assmann, D. S. Johnson, D. J. Kleitman, and J.-T. Leung, "On a dual version of the one-dimensional bin packing problem," *Journal of algorithms*, vol. 5, no. 4, pp. 502–525, 1984.

[11] J. G. Shanthikumar and D. D. Yao, "Multiclass queueing systems: Polymatroidal structure and optimal scheduling control," *Operations Research*, vol. 40, no. 3-Supplement-2, pp. S293–S299, 1992.

[12] S. Bikhchandani, S. de Vries, J. Schummer, and R. V. Vohra, "An ascending Vickrey auction for selling bases of a matroid," *Operations research*, vol. 59, no. 2, pp. 400–413, 2011.