Towards “A Theory of Spectrum Zoning”

Anant Sahai, Kristen Woyach, Kate Harrison, Hari Palaiyanur, and Rahul Tandra
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley, CA
Emails: {sahai, kwoyach, harriska, hpalaiya, tandra}@eecs.berkeley.edu

***Abstract*— Spectrum zoning is the problem faced by regulators: how to allocate different bands to different rules for radio operation. This problem has largely fallen through the cracks since to the extent that they even recognize the problem’s existence, traditional communities have viewed this as “somebody else’s problem.” Engineers/information-theorists assume that economists/lawyers will handle this and economists/lawyers put their trust in some varying combination of engineers and “the market” viewed as a pseudo-religious entity.**

This preliminary investigation puts forth a semi-Rawlsian perspective on zoning: viewing it as a robust optimization problem ($\max\min\max$ for now) in which zoning decisions are made at a slow time-scale in the face of both natural variation in the suitability of different bands for different uses *and significant uncertainty regarding user preferences*. Consequently, it can make sense to choose band-plans that are not on the Pareto frontier: but only if their runtime flexibility is more valuable than the performance overhead they impose. We examine the issues in the context of a toy model of TVs and wireless ISPs.

I. Introduction and Core Issues

[1]Spectrum zoning is the quintessential S.E.P. in the area of wireless communication. The useful S.E.P. concept was most memorably put forth by Douglas Adams in [1]:

> “An S.E.P. is something we can’t see, or don’t see, or our brain doesn’t let us see, because we think that it’s somebody else’s problem. That’s what S.E.P. means. Somebody Else’s Problem. The brain just edits it out, it’s like a blind spot. If you look at it directly you won’t see it unless you know precisely what it is. . . . ”
>
> They could now clearly see the ship for what it was simply because they knew it was there. It was quite apparent, however, that nobody else could. . . . The Somebody Else’s Problem field . . . relies on people’s natural predisposition not to see anything they don’t want to, weren’t expecting or can’t explain.

Traditionally in wireless communication, the goal is to maximize performance subject to certain implementation constraints. Where these constraints come from is sometimes discussed, but for many wireless practitioners, coming up with the constraints (like which frequency to use or the allowed transmit power) is Somebody Else’s Problem.

The emerging area of cognitive radio challenges some of the constraints. A typical cognitive radio talk starts with the NTIA Spectrum Allocation chart of Figure 1 and then

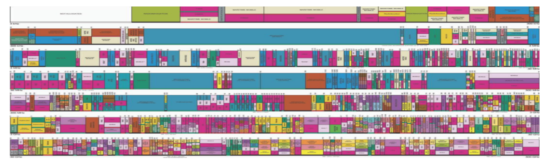

Fig. 1. What spectrum zoning looks like today. Each band has its own rules of operation for devices and in most cases, an intended use. The actual assignment of users within these bands is handled by separate processes that sometimes involve auctions, but the zoning itself is done by regulators.

argues that in practice, the spectrum appears to be vastly underused (citing for example [2]). This inefficiency has serious economic consequences. As early as 1992, the social welfare surplus of reallocating just the UHF TV bands to cellular was estimated at about \$1 billion [3]. However, the S.E.P. effect is so strong that even in the cognitive radio community, the core challenge of spectrum zoning — *how to assign regulatory rules governing the devices authorized to operate in various bands* — gets evaded in favor of issues related to the assignment and use of already-zoned spectrum. These include algorithmic questions like how to sense underutilized spectrum, how to run a real-time spectrum market, how to cooperate to deliver messages, etc. Zoning itself remains Somebody Else’s Problem.

A. Current thinking

To the extent that it occurs at all, discussion of zoning tends to be centered on the economics/law side. Although there is broad agreement that ideally there should be fewer zones than currently exist today, spectrum policy advocates tend to fall into three different camps: believers in expanded spectrum commons where the restrictions are on the equipment power alone, “privatizers” who want usage rights to be traded like property, and those who believe in the mixed “cognitive radio” approach to allow different priority classes of users to coexist in the same band. We focus here on the latter two camps because wildly heterogeneous use is harder to support in a pure commons due to the one-size-fits-all requirements. Faulhaber in [4] calls this the power-mix problem.

The underlying problem is believed to be a lack of flexibility that would allow mismatches to be corrected. Markets are often proposed as the right mechanism to provide the requisite flexibility [5], [6]. Would the kind of ‘big-bang’ auction proposed in [7] be sufficient to allocate the bands correctly? We know the auction itself must be designed well [8], [9]. But, even with a perfect auction, how do we define the parcels we are auctioning off? Are all parcels equal?

[1]This invited paper admittedly represents work in progress. As a result, the scholarly treatment of existing perspectives and prior work here is not quite up to our own standards and the technical story is not quite as tight as we would like it to be either. Those interested are directed to see our submissions to DySpAN ’10 for a fuller and more polished account.

Hatfield and Weiser suggest that parcels of spectrum are not in fact equal, even from a regulators perspective [10]. At the very least, the different propagation characteristics require different approaches to enforcement. For example, AM radio band propagation varies greatly with the state of the ionosphere, so trying to maintain sharp demarcation lines between different users in these bands is futile. This kind of propagation also lends itself most naturally to broadcast systems, not short-range unicast systems. (This will also show up in our own toy model later.) PCS bands, by contrast, are much more predictable. This suggests that a nontrivial zoning of spectrum is likely necessary [11].

But how should we decide what should be possible within a given band? Perhaps auctions can be used here too in order to decide the usage rules for the band depending on the current (and forecast) demand [12]–[14]. Unfortunately, Hatfield and Weiser issue a warning on this issue as well – it is entirely possible that the usage rules for a band may in themselves preclude the 'best' uses [15], and we might never even know this had happened.

The warnings of Hatfield and Weiser seem to be largely unheeded by the spectrum policy community. Hazlett's response [16] is representative: the courts and markets are strong enough to correct any mistakes made in the original planning. Courts can soften and modify the rules as practiced and market forces can allow aggregation of spectrum into critical masses held by private operators that can just internally change how they enforce policies, provided that they have enough freedom of use to do so. They can in turn lease it out again with more suitable rules. Either way, property rights will evolve toward their optimal forms. However, it is not clear that this statement actually has technical justification, or *even what a technical justification for this would look like.* In other words, the details are Somebody Else's Problem.

The underlying core issue is one of time-scales and dynamic change. As Hazlett rightly points out, there is a need for rules to be stable long enough for both engineers to *innovate* systems in response to them *and* for investors to recoup their fixed costs. So nobody expects either courts or administrative bodies to change rules too rapidly. Furthermore, there is no reason to believe that the existing courts and system of common law can react to technological developments and variations in user preferences any faster than an administrative body like the FCC. If anything, their respect for tradition and the legal requirement to act in the face of actual controversy make common-law converge even more slowly than administrative processes.

This is a challenge in the context of radio systems that can ride Moore's law. There is also an engineering dimension to the problem that is different: we cannot engineer humans to have fundamentally new characteristics (at least, not yet). So the law must take people as we find them. Radios, on the other hand, can be mandated to have a wide range of possible features and this is what the rules can do.[2]

[2]This is a different argument than the one between specialized vs general courts and the value of technical expertise among adjudicators.

For property rights to evolve in practice, the rules governing radios must be flexible enough to allow different users and different systems to trade spectrum while *maintaining enforceability.* Light-handedness is a prerequisite for flexibility, and while light-handed mechanisms exist to enable enforceability (e.g. to avoid the "hit and run radio" problem identified by Faulhaber in [4], [17]), they come with a nontrivial[3] cost [18]–[20]. When is this cost worth paying? There are genuine technical questions that cannot simply be swept under the rug when considering how to restructure spectrum.

B. Our broad approach

This paper takes some baby steps in trying to understand what is actually important to consider in spectrum zoning and band allocation. A toy model is introduced in Section II that simplifies the world as having just two kinds of radios: broadcast television and wireless Internet service providers (ISP) (e.g. WiMax, 802.22, etc.). This allows the issues of propagation characteristics and heterogeneous coexistence to be explored both in an extremely idealized toy simulation with a single TV tower and with respect to the real-world distribution of TV towers in the USA. Once the utility model is clear, Section III allows us to illustrate the Pareto frontier to see how whitespace sharing enables us to push that out.

Finally, Section IV is where we discuss how regulators should choose a band-plan. It is here that our perspective is inspired by John Rawls' "Theory of Justice" in which he explores how rules should be set for human political society. His perspective of "Justice as Fairness" uses a thought-experiment he calls "the original position" — the hypothetical perspective of unborn people who have not yet been dealt a particular life and must decide on the rules under which they are going to live. Under this veil of ignorance, Rawls claims they should adopt two principles: [21]

1) Each person has an equal right to "fully adequate" flexibility and autonomy (basic liberties) in a scheme that is compatible with the same liberties for all.
2) Any inequalities that develop must be of the greatest benefit to the least-advantaged members of society.[4]

Rawls' original position is an extreme version of that in which regulators find themselves: they cannot know what user preferences will be and what technological innovation will bring. The political rules of society correspond to a spectrum zoning decision. We will adopt the $\max\min$ perspective of Rawls' "difference principle" (point 2 above). The main distinction is that unlike Rawls, we will reverse the priority of these two points. For us, flexibility is not a non-negotiable requirement since we cannot assume that citizens are willing to radically sacrifice utility for the sake of radio flexibility

[3]This is different from "transactions costs" that can be amortized away as you use the resource. Instead, the mechanisms to preserve enforceability impose ongoing obligations that inevitably reduce the productive value of the resource. They are like a "tax" that is paid in-kind.

[4]Simplified from the exposition in [21] where this second point also contains the requirement that the inequalities/benefits must be attached to positions that are open to all with fair equality of opportunity. While this omitted point does figure in related discussions of wired network neutrality, it does not seem relevant to the problem of wireless spectrum zoning.

and autonomy.[5] As a result, flexibility turns out to be the right choice only when it does not cost too much. Mathematically, this is seen by formulating the problem as a $\max\min\max$ problem.[6] The most interesting feature of this formulation is that it tells us that operating on the Pareto frontier need not be the right choice for a band plan.

II. A Toy Model for Spectrum Zoning

Suppose the spectrum consists of a contiguous band located between f_l Hz and f_u Hz, divided into N_c channels, each of bandwidth $W_c = (f_u - f_l)/N_c$ Hz. For our simulations, we take $f_l = 50$MHz and $f_u = 350$MHz, for 50 channels of width 6MHz each. This roughly corresponds to the VHF and UHF digital TV bands in the USA[7].

Now, any particular channel can be allocated for one of three rules: ISP, TV or 'Shared.' Shared means a potential 'white space' [22] channel in which TV transmission is the primary use, but secondary ISP service can be provided as long as it does not harm the primary within its protected area [23]. Hence a 'band plan' is a function F that designates the use of each channel in the band as well as designates a measure of protection for the primary in those channels that are shared. In our case, the protection is parametrized by a single non-negative number, or 'erosion margin', as described in [23], so a band plan is of the form

$$F : \{1, 2, \ldots, N_c\} \to \{\text{TV}, \text{ISP}, \text{Shared}\} \times \mathbb{R}.$$

For simulations, we additionally restrict band plans to have uses allocated contiguously and a single erosion parameter. Hence, there are 6 possible orderings, e.g. [TV, Shared, ISP], [ISP, Shared, TV], [Shared, ISP, TV], etc. Figure 2 shows an example of such a simplified band plan. Let $\mathcal{F}$ be the set of all such contiguous band plans.

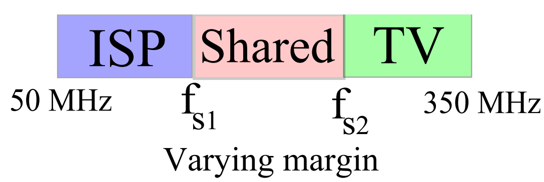

Fig. 2. An example band plan. The shared channels have a variable erosion margin parameter that changes the protected area of the TV primary, but within any given band plan, only one value for this margin is used.

A. Propagation characteristics of channels

We assume that transmitters for TV and ISP transmit at powers $P_{TV} = 10^6$ and $P_{ISP} = 4$ Watts per channel respectively. The wireless medium is modeled as an additive white noise Gaussian noise (AWGN) channel with a frequency-dependent path loss and (room temperature) thermal noise level N_0 (dBW/Hz). The received power of a signal at distance r away in a channel with center frequency f_c is modeled using the generalized Friis transmission equations [24] as

$$P_{TV,r}(f_c) = P_{TV} \left(\frac{1}{r}\right)^{d_{TV}} \left(\frac{c}{4\pi f_c}\right)^2,$$

where c is the speed of light and d_{TV} is the TV path-loss exponent. For an ISP signal, d_{ISP} is the generally larger ISP path-loss exponent.

Within a TV channel, it is assumed that a receiver can properly decode the TV signal if and only if its SINR is above a certain operational limit, Δ_{TV}.

B. TV Coverage

We define the noise-limited radius of the TV tower in a channel as the distance at which the receiver's SNR drops to Δ_{TV}. In this model,

$$r_{nl}(f_c) = \left(P_{TV} \left(\frac{c}{4\pi f_c}\right)^2 (\Delta_{TV} N)^{-1}\right)^{1/d_{TV}},$$

where $N = W_C 10^{(N_0/10)}$ Watts is the channel noise power.

Figure 3 shows what our toy world looks like in one particular channel if the channel is allocated to either exclusive TV use or exclusive ISP. In the case of an exclusively TV channel, the coverage of the TV extends out to a radius of $r_{nl}(f_c)$. We assume the distance between TV transmitters, denoted as $2R_{MAX}$ in Figure 3, is large enough so that the interference from other TV transmitters is on the order of the noise power at the noise-limited radius in our world.

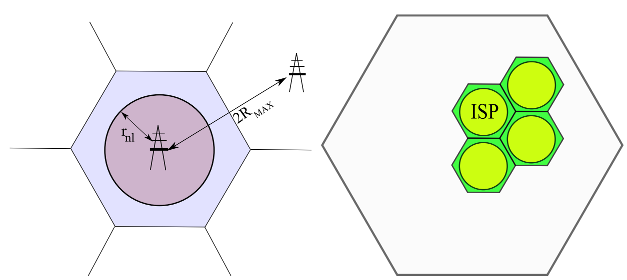

Fig. 3. Allocation of a particular channel within the band to one of two uses: TV or ISP exclusively. The hexagonal division tiling of area is done for visualization. In simulations, we assume our world consists of one circle of radius R_{MAX}.

Before moving on to a discussion of how we measure the utility of band plans, we briefly discuss the notion of an erosion margin in shared channel where TV transmitters are primary users. Figure 4 our toy world in a channel that is shared. The tradeoff between secondary use with primary use in a white-space model, is adjusted by varying the erosion margin $\gamma \geq 0$ (dB) [22], [23]. The margin is an erosion in the SINR that must be accepted by the primary receivers, which leads to a reduction in the service area. The *protected radius* $r_p(\gamma, f_c) < r_{nl}(f_c)$ is set as the radius

[5]Although the first author is writing this footnote using GNU/Emacs and LaTeX on a system running GNU/Linux, it is likely that many of you are reading this on a system that is not mostly Free Software. The priority of liberty over all other factors cannot be taken as a given among the citizenry when it comes to technological matters.

[6]The inner $\max$ is also implicit in Rawls' account, but it is not discussed explicitly because Rawls himself wants to treat the $\max\min$ formulation as a way to capture other politically desirable concepts like reciprocity rather than as a self-evident way to proceed [21].

[7]In reality, TV channels are not allocated entirely contiguously.

at which the TV signal SNR (in dB) falls to $\Delta_{TV,dB} + \gamma$, where $\Delta_{TV,dB} = 10\log_{10}\Delta_{TV}$. In addition, a *no-talk* radius $r_n(\gamma, f_c)$ is prescribed so that interference from allowed secondary use (beyond the no-talk radius) keeps the the primary TV signal's SINR at $r_p(\gamma, f_c)$ at at least Δ_{TV}. There are several ways of calculating $r_n(\gamma, f_c)$ based on the model for interference, see Chapter 4 of [25].

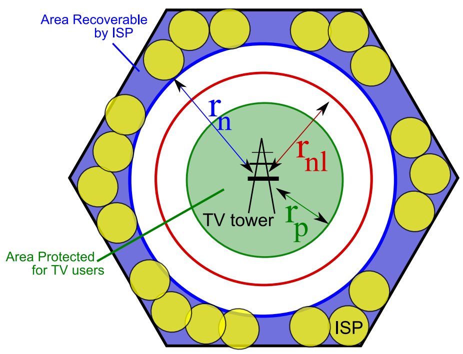

Fig. 4. Shared use of a channel where the TV system is the primary user and ISPs are secondary users.

C. Utility calculations for TV and ISP

When the channel is designated for unshared or shared TV use, we ascribe a TV utility for the channel equal to its guaranteed coverage area. In the unshared case, this is $\pi r_{nl}(f_c)^2$, while in the shared case, it is $\pi r_p(\gamma, f_c)^2$.

Unlike TV's natural multicast nature, the ISP is a series of unicast flows. So, it is not sufficient to just consider the area served, we must instead consider the capacity an individual user would get over time, which we abstract into the capacity the ISP could support per unit area. Each user sees interference from its neighbors. Therefore, we must consider an exclusion radius around each ISP user in which the others are not allowed to transmit [26]. The exclusion radius serves as a proxy for the MAC protocol that coordinates the sharing of resources amongst ISP users.

For the purposes of our toy model, the interference is assumed to be coming from six interferers located at the boundary of the exclusion radius. If the exclusion radius is large, the interference is low, but the capacity per unit area takes a hit from lower frequency reuse. The capacity per unit area is optimized over the exclusion radius. No capacity is available where the ISP is not permitted to operate.

We also consider two different types of ISP: short range (.25km from the ISP's tower) and long range (2.5km from the ISP's tower). These exhibit significantly different behavior in terms of how they prefer different frequencies and how they react to pollution from a TV primary. Figure 5 considers the utility-ratio between an equally sized channel at high and low frequency. We see markedly different behavior. TV always prefers the lower frequencies: this verifies our intuition that low frequencies have better propagation characteristics and so the TV can cover more area in these bands. The short-range ISP is relatively indifferent to the location of the channel, while the long-range ISP prefers the lower frequencies.

Parameters	Description	Values
f_l	Lower band limit	50 MHz
f_u	Upper band limit	350 MHz
W_c	Channel bandwidth	6 MHz
N_c	Number of channels	50
N_0	Noise power	-204 dBW/Hz
d_{TV}	TV path loss exponent	3.5
d_{ISP}	ISP path loss exponent	4
c	Speed of light	3×10^8 m/s
R_{MAX}	Radius of toy world	200 km
γ	Margin (primary erosion)	0 to 10 dB
ISP Range	Radius of ISP 'cell'	.25 km, 2.5 km
P_{TV}	TV transmitter power	1 MW
P_{ISP}	ISP transmitter power	4 W
P_{INT}	ISP interferer power	4 W
N_{INT}	Number of ISP interferers	6
$\Delta_{TV,dB}$	TV SNR cutoff	18 dB

TABLE I
VALUES OF SIMULATION PARAMETERS.

When the ISP tower is very close to the user, the SNR at that user is extremely high: the capacity is degree-of-freedom limited. The short-range ISP cares only that it is getting 6MHz, not where this 6MHz happens to be located. The long-range ISP, on the other hand, has a lower SNR and is power limited. So, the transmission characteristics do matter and the long-range ISP prefers lower frequencies.

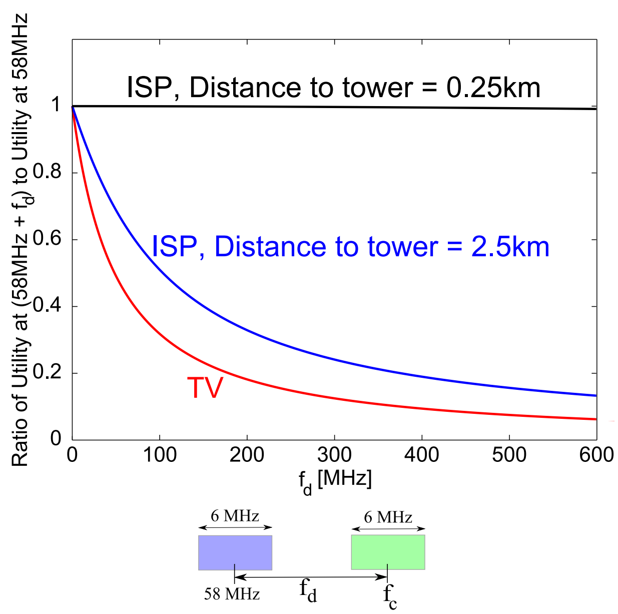

Fig. 5. Ratio of the utility gained in a 6MHz band at a higher frequency to that gained in the same bandwidth at 58MHz.

The difference in bandwidth-limited and power-limited can also be seen in the reaction of the two ISPs to pollution[8] from a primary TV tower. This effect is shown in Figure 6. The plot shows the ratio of utility gained in a polluted 6MHz channel

[8]In shared bands, the local pollution level coming from TV signals also figures significantly into the optimization of the MAC radius. This is accounted for in the simulations. So while the erosion margin for the primary users is set once for the whole systems, the MAC radii for ISP secondaries vary dynamically from place to place.

to the utility gained from a clean one. The large-scale ISP is power-limited and therefore sensitive to pollution that limits its SNR. The short-range ISP is bandwidth-limited, so it does not care much about small amounts of pollution. Of course, as the amount of pollution gets very big, the short-range ISP begins to suffer as well.

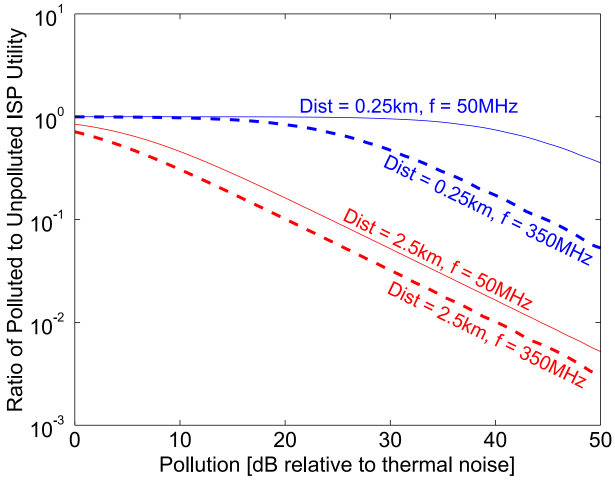

Fig. 6. Ratio of the utility of a polluted channel to the utility in a clean channel. The long-range ISP is much more sensitive to pollution, so it is less likely to be able to pick up extra capacity from sharing bands with TV stations. Short-range ISPs, on the other hand, do not mind the pollution as much and so can recover white spaces around the TV protection area.

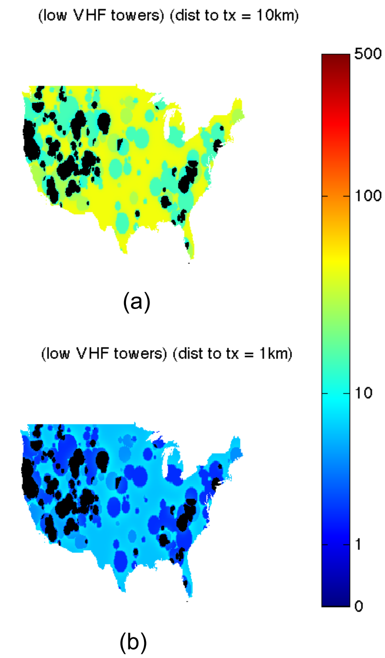

Fig. 7. Ratio of the capacity gained in the polluted low VHF channels 5,6,7 (18 MHz total) to that gained in a clean 700MHz channel and expressed in terms of effective MHz. The long-range ISP greatly prefers the low frequencies while the short-range ISP is less sensitive to changes in frequency and sees a much lower ratio. In these maps, the ISPs must respect the protected contours of the TV stations as outlined in [27].

The difference between the ISPs can also be seen by taking the same model for the ISP exclusion region and adapting[9] it to real data from the FCC. The TV towers are placed in their actual locations on the U.S. map. We then assume that ISPs must respect the no-talk radius imposed by the protection radius rules outlined in the FCC Whitespace ruling [27]. These rules stipulate that the co-channel no-talk radius is 14.4km outside of the protection radius, and the adjacent channel no-talk radius is 0.74km outside of the protected radius. In our model, the ISPs treat the pollution caused by the TV stations as noise. We then normalize this capacity per unit area in the polluted TV bands by dividing through by the capacity per unit area in a clean 700MHz band.

Figure 7 shows just the low VHF channels: 5, 6, 7. The long-range (10km) ISPs in most of the country greatly prefer even the mildly polluted TV bands to the bad propagation characteristics at 700MHz. The short-range (1km) ISP finds the 700MHz bands very good, and the polluted TV bands almost equivalent. The interstices around the TV protected areas are therefore quite attractive to the shorter-range ISPs.

Figure 8 shows this same ratio of capacity in the polluted TV bands to capacity in a clean 700MHz band for all TV channels without imposing the FCC imposed protection radius. The TV towers are still producing pollution, but the government is not barring the ISPs from any locations. With all the towers included, we can see the effect of the short range ISP (top left corner) getting greater utility out of the polluted channels than the longer range ISP (bottom right corner). Figure 9 shows this same calculation and includes the protection radius around the TV imposed by the FCC[10]. The interesting point here is the difference between this and the previous Figure: the limiting factor is the FCC requirements on protection radius for TV. This suggests that self-interest (avoiding polluted bands) will not keep ISP users out, hence rules for sharing are required [4]. This effect is particularly significant for short-range ISPs.

III. The Pareto Frontier

Given a band plan F, let $U_{TV}(F)$ denote the total coverage area of TV transmission summed over channels allocated to TV, and let $U_{ISP}(F)$ denote the ISP-capacity summed over channels and areas where the ISP can legally operate. Define $\mathcal{U} =: \{(U_{TV}(F), U_{ISP}(F)) : F \in \mathcal{F}\}$ be the set of utility points achieved by all band plans. A given point (u^P_{TV}, u^P_{ISP}) belongs to the Pareto frontier $\mathcal{P}$ if there does not exist any point $(u_{TV}, u_{ISP}) \in \mathcal{U}$ such that both $u_{TV} > u^P_{TV}$ and $u_{ISP} > u^P_{ISP}$. The associated band plan is Pareto-efficient.

The first row of Figure 10 shows the Pareto frontier for two ISP ranges. Note that this plot shows that different size ISPs prefer sharing to different extents. Short range ISPs in this model are degree of freedom limited, so they would prefer the bandwidth if it is available, even if it is polluted by primary TV use. Long range ISPs however are SNR

[9]The FCC's recommended propagation models are used as well as the actual transmission powers for the television towers involved.

[10]Including the need to protect TV receivers tuned to adjacent channels. In fact, this is what is dominating the story in many places. For more on all this, see [28].

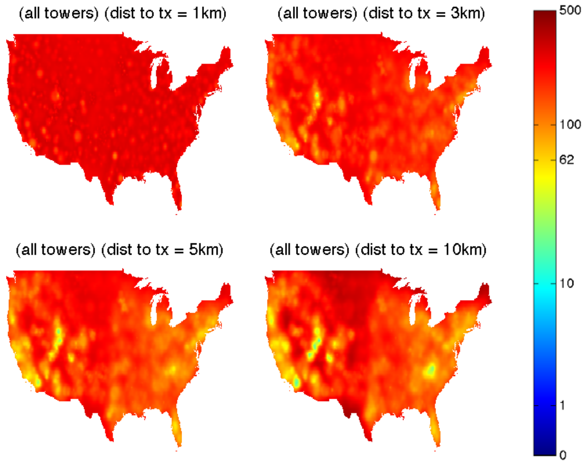

Fig. 8. Ratio of capacity per unit area in all of the (polluted) TV bands to the capacity in a clean 700MHz band. There is some difference between the ISPs because the long-range ISPs are power-limited and are therefore much more sensitive to pollution. But there is a lot of available capacity here when the TV pollution is treated as noise and the ISPs do not have to respect the TV protection rules defined by the FCC.

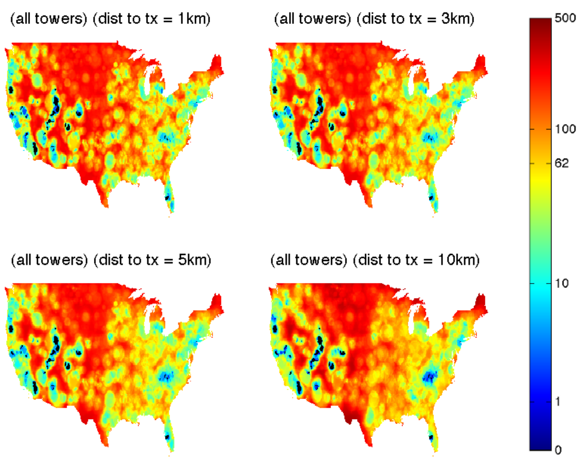

Fig. 9. The same ratio as in previous plots, but with all TV towers and with ISPs respecting the protected contour defined by the FCC. The protection contour is the limiting factor as these plots display far less capacity available than Figure 8

limited, so beyond some point, they would prefer to just have unpolluted exclusive ISP bands.

However, coexisting with primaries costs secondary users some amount of overhead due to hardware complexity [29], sensing [30], enforcement [31], [19] and protection [32]. Now, overhead can be viewed as cutting the secondary ISP utility in shared bands, but leaving the unshared bands the same. The second row of Figure 10 shows what happens when half the utility in shared bands is cut out and the third row shows what happens when 95% of the utility in the shared bands is cut out.

IV. Spectrum Zoning as Robust Optimization

Now that we have the means of measuring utility for the two uses of spectrum, we discuss a straightforward way of

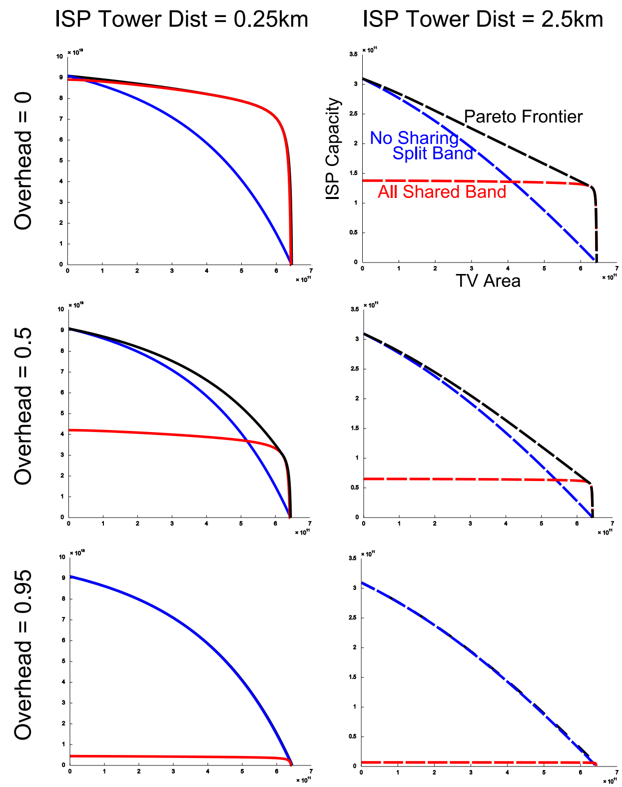

Fig. 10. Pareto frontiers for two ISP ranges and three values of overhead. Red lines show points hit by fully sharing the spectrum and varying the erosion margin, blue lines denote splitting the spectrum between exclusive TV and ISP use and black lines show the Pareto frontier amongst all band plans. Note that in the .25km ISP with 0 overhead, fully sharing the band essentially achieves the Pareto frontier, while both ISPs with high (0.95) overhead, splitting the band for exclusive use essentially achieves the Pareto frontier.

how a regulatory agency might choose a band plan. Let $\rho \in [0,1]$ denote the user preference for TV and $1-\rho$ for ISP. The utility of a band plan is then

$$U(F,\rho) = \rho U_{TV}(F) + (1-\rho) U_{ISP}(F). \tag{1}$$

If we knew the preference parameter ρ, the best band plan would be the point of the Pareto frontier that has a slope of $\frac{-\rho}{1-\rho}$. See Figure 11 (a).

A. *Uncertain preferences*

We now consider the case of uncertain ρ. As the time-scale at which zoning decisions are made is much larger than the time-scale of user preferences, it is unreasonable to assume any *a priori* probabilistic model for ρ. Instead, we assume that ρ can take any value within a given interval $[\rho_{min}, \rho_{max}]$. Given this uncertainty model, it is reasonable to maximize the worst case utility:

$$U = \max_{(u_{TV}(F), u_{ISP}(F)) \in \mathcal{U}} \min_{\rho \in [\rho_{min}, \rho_{max}]} U(F,\rho). \tag{2}$$

Since the utility function is linear in ρ it is easy to see that

$$\min_{\rho\in[\rho_{min},\rho_{max}]} U(F,\rho) = \begin{cases} U(F,\rho_{min}) & \text{if } u_{TV}(F) > u_{ISP}(F) \\ U(F,\rho_{max}) & \text{if } u_{TV}(F) \le u_{ISP}(F) \end{cases}$$

Intuitively (as Figure 11 shows), if $u_{TV} \le u_{ISP}$ then the worst-case uncertainty corresponds to the least possible fraction of people liking TV, i.e., $\rho = \rho_{min}$. So, the max-min zoning decision is the reasonable Pareto point that is closest to the point where the line $u_{TV} = u_{ISP}$ intersects the frontier.

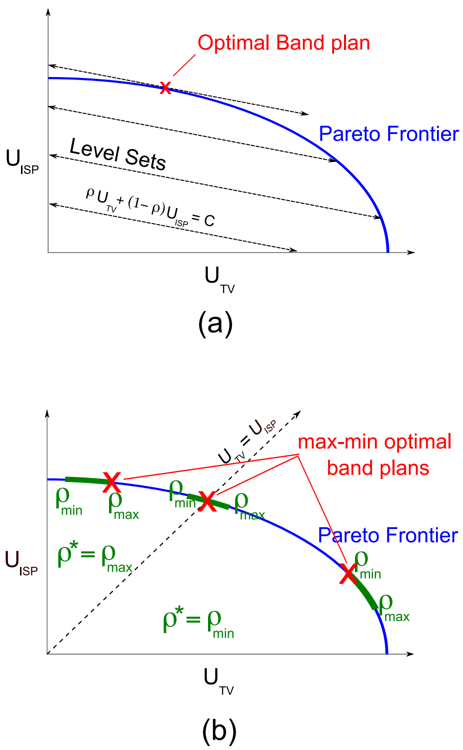

Fig. 11. The max-min optimization framework to evaluate the zoning decision. Part (a) shows the optimal zoning decision for a fixed ρ. Part (b) illustrates the case when ρ is uncertain.

B. Flexible band plans

The common thread between currently proposed spectrum allocation approaches is some level of flexibility on a time scale smaller than the rule-setting time scale. However, there will be some aspects of the band plan that cannot change at runtime. To model this we assume that given a band plan F, there exists a run-time adjustable parameter κ (e.g. assignment of bands) that allows choice between multiple points in the utility space. Hence, the optimization problem is

$$U = \max_{F\in\mathcal{F}} \min_{\rho\in[\rho_{min},\rho_{max}]} \max_{\kappa} U(F,\rho,\kappa), \tag{3}$$

where $U(F,\rho,\kappa)$ is the total utility for a given band plan F, user preference ρ and the run-time parameter κ chosen in response to the user preference ρ.

Figure 12 illustrates the max-min-max optimal band plan for two different flexible band plan scenarios. When the flexible band plan operates close to the Pareto frontier, the max-min-max optimizer chooses the flexible band plan over any band plan on the Pareto frontier. However, the flexible band plan is ineffective if it operates too far away from the Pareto frontier. The level of overhead matters.

As discussed in Section III, flexibility and sharing incur certain overheads to maintain peaceful coexistence of systems that might use the same band [19], [29]–[32]. These overheads[11] would need to be set at regulation time, not at runtime. Likewise, the sharing protocol, such as the erosion margin to determine r_p for the shared bands would need to be set beforehand. At runtime you could only change whether the shared bands will be actually used for shared spectrum or whether an ISP will buy out a TV station's primary rights and the band will effectively convert to a clean ISP band. But since the overheads must be set at regulation time, such an ISP would still incur the overheads involved.

Figure 13 shows the utility points achieved by flexible band plans when the primary (TV) also incurs an overhead when operating in a shared channel in terms of a fraction of power that is diverted by overheads. Note that an individual band plan with many shared channels can achieve different utility points at runtime. When the margin is small, as seen in the left column of the figure, all shared channels in which the TV primary is operating are essentially given entirely to TV, while ISP recovers no utility in these channels. When the channel is shared, but allocated at runtime to ISP, it achieves the same utility it would in an exclusive channel if the overhead is 0. As overhead increases, this utility is scaled down and it becomes unclear whether flexibility buys much for the overall system in terms of robust optimization.

Acknowledgments

We thank the United States National Science Foundation (CNS-403427, CCF-729122, CNS-0932410 and a Graduate Research Fellowship) for their generous funding support. Mubaraq Mishra and Pulkit Grover are thanked for very useful discussions.

References

[1] D. Adams, *Life, the universe, and everything (Hitchhiker's guide to the universe trilogy)*. Del Ray, 1983.

[2] M. A. McHenry and K. Steadman, "Spectrum occupancy measurements, location 1 of 6: Riverbend park, Great Falls, Virginia," Shared Spectrum Company, Tech. Rep., 2005. [Online]. Available: http://www.sharedspectrum.com/inc/content/measurements/nsf/1_NSF_Riverbend_Park_Report.pdf

[3] E. R. Kwerel and J. R. Williams, "Changing Channels: Voluntary Reallocation of UHF Television Spectrum," *FCC Office of Strategic Planning Working Paper Series*, no. 27, Nov 1992.

[11] Principal amongst necessary overheads incurred because of flexibility are those that show up at device manufacture. In addition to hardware complexity (which can be amortized over the use), there is increased power-consumption (which is paid on an ongoing basis). More significantly (since these are not subject to Moore's law), there are the resource-costs associated with any enforcement-related requirements that must be certified. For example, identity codes for enforcement and the sensing margins set for protection of TV service areas [20]. Certification occurs once.

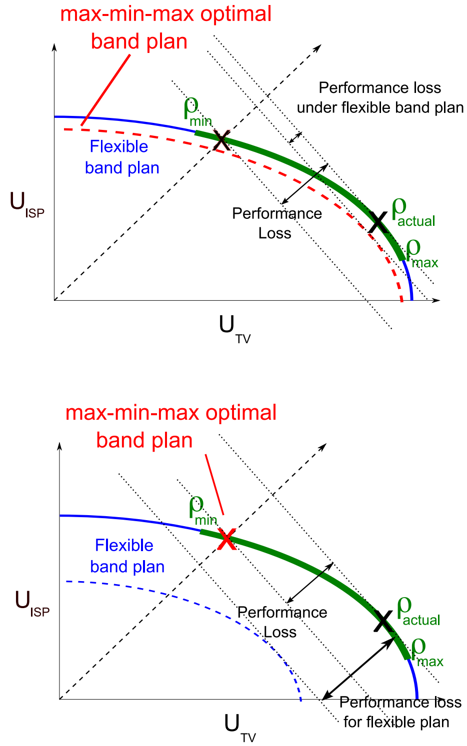

Fig. 12. Flexible band plans can achieve multiple points in the utility space at run-time after the zoning has occurred. If overheads are too high, non-flexible plans that operate nearer to the Pareto frontier would be better even with run-time uncertainty.

[4] G. R. Faulhaber, “Deploying Cognitive Radio: Economic, Legal, and Policy Issues,” *International Journal of Communication*, vol. 2, pp. 1114–1124, 2008.

[5] R. H. Coase, “The Federal Communications Commission,” *Journal of Law and Economics*, vol. 2, pp. 1–40, Oct. 1959.

[6] A. S. de Vany, R. D. Eckert, C. J. Meyers, D. J. O’Hara, and R. C. Scott, “A Property System for Market Allocation of the Electromagnetic Spectrum: A Legal-Economic-Engineering Study,” *Stanford Law Review*, vol. 21, no. 6, pp. 1499–1561, Jun. 1969.

[7] E. Kwerel and J. Williams, “A Proposal for a Rapid Transition to Market Allocation of Spectrum,” *FCC Office of Strategic Planning Working Paper Series*, no. 38, Nov 2002.

[8] J. McMillan, “Why Auction Spectrum,” *Telecommunications Policy*, vol. 19, no. 3, pp. 191–199, Apr 1995.

[9] ——, “Selling Spectrum Rights,” *Journal of Economic Perspectives*, vol. 8, no. 3, pp. 145–162, Summer 1994.

[10] D. Hatfield and P. Weiser, “Toward Property Rights in Spectrum: The Difficult Policy Choices Ahead,” *CATO Institute*, no. 575, Aug. 2006.

[11] P. J. Weiser and D. N. Hatfield, “Policing the Spectrum Commons,” *Fordham Law Review*, vol. 74, no. 2, pp. 663–694, 2005.

[12] M. M. Bykowsky, K. R. Carter, M. A. Olson, and W. W. Sharkey, “Enhancing Spectrum’s Value Through Market-informed Congestion Etiquettes,” *FCC Office of Strategic Planning Working Paper Series*, no. 41, Feb 2008.

[13] M. M. Bykowsky, M. A. Olson, and W. W. Sharkey, “Modeling the Efficiency of Spectrum Designated to Licensed Service and Unlicensed Operations,” *FCC Office of Strategic Planning Working Paper Series*, no. 42, Feb 2008.

[14] ——, “A Market-based Approach to Establishing Licensing Rules: Licensed Versus Unlicensed Use of Spectrum,” *FCC Office of Strategic Planning Working Paper Series*, no. 43, Feb 2008.

[15] P. J. Weiser and D. Hatfield, “Spectrum Policy Reform and the Next Frontier of Property Rights,” *George Mason Law Review*, vol. 60, no. 3, 2008.

[16] T. W. Hazlett, “A Law and Economics Approach to Spectrum Property Rights: A Response to Weiser and Hatfield,” *George Mason Law Review*, vol. 15, no. 4, Summer 2008.

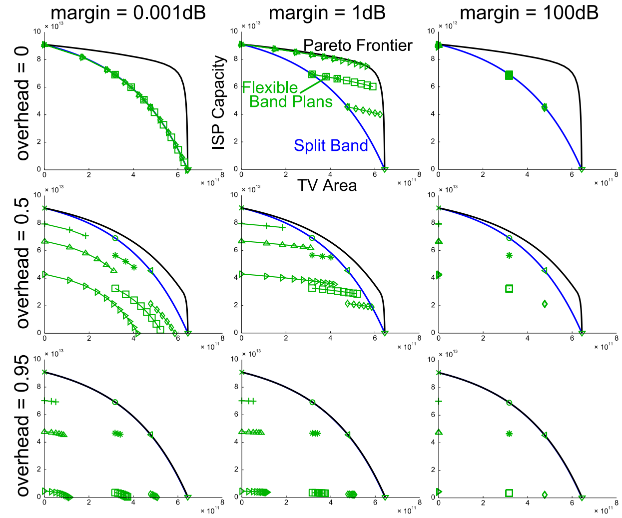

Fig. 13. Flexible band plans for the 0.25km ISP case. Three different margins and overheads are looked at in separate subfigures. Within each subfigure, flexible band plans are plotted with the specified margin under the specified overhead. The connected points are individual band plans with runtime flexibility.

[17] G. R. Faulhaber, “Wireless telecommunications: Spectrum as a critical resource,” *Southern California Law Review*, vol. 79, no. 3, Mar. 2006.

[18] G. Atia, A. Sahai, and V. Saligrama, “Spectrum enforcement and liability assignment in cognitive radio systems,” in *Proceedings of the Third IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Chicago, IL, Oct. 2008.

[19] A. Sahai, K. Woyach, G. Atia, and V. Saligrama, “A technical perspective on light-handed regulation for cognitive radios,” *IEEE Communications Magazine*, pp. 96–102, Mar. 2009.

[20] A. Sahai, S. M. Mishra, R. Tandra, and K. A. Woyach, “Cognitive radios for spectrum sharing,” *IEEE Signal Processing Magazine*, Jan. 2009.

[21] J. Rawls, *Justice as Fairness: A Restatement*, 2nd ed., E. Kelly, Ed. Belknap Press of Harvard University Press, 2001.

[22] R. Tandra, S. M. Mishra, and A. Sahai, “What is a spectrum hole and what does it take to recognize one?” *Proceedings of the IEEE*, Jan. 2009.

[23] M. Mishra and A. Sahai, “How much white space is there?” EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2009-3, Jan. 2009. [Online]. Available: http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-3.pdf

[24] T. S. Rappaport and T. Rappaport, *Wireless Communications: Principles and Practice (2nd Edition)*. Prentice Hall PTR, 2001.

[25] N. Hoven, “On the feasability of cognitive radio,” Master’s thesis, University of California, Berkeley, Berkeley, California, May 2005.

[26] J. A. Fuemmeler, N. H. Vaidya, and V. V. Veeravalli, “Selecting transmit powers and carrier sense thresholds for CSMA protocols,” University of Illinois at Urbana-Champaign, IL, Tech. Rep., Oct. 2004.

[27] “FCC Adopts Rules for Unlicensed Use of Television White Spaces.” http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-260A1.pdf.

[28] K. Harrison, S. M. Mishra, and A. Sahai, “How much white-space capacity is there?” in *IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, In preparation for, Mar. 2010.

[29] D. Cabric, S. M. Mishra, and R. W. Brodersen, “Implementation issues in spectrum sensing for cognitive radios,” in *Asilomar Conference on Signals, Systems, and Computers*, 2004.

[30] R. Tandra and A. Sahai, “SNR walls for signal detection,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 1, pp. 4–17, Feb. 2008.

[31] K. A. Woyach, A. Sahai, G. Atia, and V. Saligrama, “Crime and punishment for cognitive radios,” *Fourty-Sixth Allerton Conference on Communication, Control, and Computing*, Sep. 2008.

[32] A. Sahai, N. Hoven, and R. Tandra, “Some fundamental limits on cognitive radio,” *Forty-second Allerton Conference on Communication, Control, and Computing*, Oct. 2004.