Cognitive radios in the TV whitespaces: challenges and opportunities

Kate Harrison

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2011-151
http://www.eecs.berkeley.edu/Pubs/TechRpts/2011/EECS-2011-151.html

December 16, 2011

Acknowledgement

NSF graduate research fellowship

Cognitive radios in the TV whitespaces: challenges and opportunities

Kate Harrison

Contents

Abstract

In the past decade, wireless devices have become increasingly popular which in turn makes spectrum more valuable and more scarce. In an effort to help alleviate the current spectrum shortage crisis in the United States, the Federal Communications Commission (FCC) has ruled that starting in 2008 unlicensed devices are allowed to transmit wirelessly in the television bands [1, 2]. This permits unlicensed devices such as a wireless router or cordless phone to transmit in the "unused" portions of the TV bands, also known as the *TV whitespaces*. The TV whitespaces are generally regarded as a great opportunity for unlicensed devices (also known as *secondaries*) to expand beyond the crowded Industrial, Scientific and Medical (ISM) bands. We explore the magnitude of this opportunity as well as the challenges it presents.

This thesis first quantifies the opportunity to secondaries in terms of available spectrum and achievable data rates via simulations involving real-world TV assignment [3] and population data [4, 5]. Prior studies for the United States [40, 41] and the United Kingdom [43] have bounded the amount of available spectrum but fail to account for significant effects such as an increased noise level and self-interference among secondaries. These effects mean that the traditional bandwidth metric is inadequate since spectrum users are not interested purely in spectrum but rather in what they can do with it, which can be characterized by the achievable data rate. Our work incorporates this metric in order to better quantify the impact of the TV towers and associated regulatory decisions on the utility of whitespaces to secondaries. We also develop models for both point-to-point and cellular systems in order to quantify the impact of self-interference among secondaries in the TV whitespaces.

Second, using these models, we compare simple spectrum reallocation to the use of whitespaces and conclude that the latter is better if we do not wish to experience drastic changes in TV availability. This improves the work done by Mishra, et al. in [40] by more accurately modeling secondary utility.

Third, we find that the current FCC regulations [2] may be inadequate to protect TV viewers from the harmful effects of aggregate interference because the rules are made implicitly with only a single transmitter in mind. The authors of [37] reach the same conclusion regarding the safety of both FCC and ECC rules in Finland but they do not suggest a remedy. We suggest setting a maximum power density for secondaries rather than a per-device power limit to avoid this problem. Such an approach might seem difficult with sensing-only technologies, but we believe that the existence of TV whitespace databases enables this necessary fine-grain control.

Given that enforcing a power density is necessary for TV receiver protection, we then explore how this solution can be employed to improve utility as well as freedom for secondary users. The current FCC rules implicitly favor certain applications. We offer a principled way to help mitigate the tensions between two types of users, rural and urban, using an approximately-optimal algorithm to choose a power density.

Finally, we provide a toolkit which is publicly available at [6]. This Matlab-based toolkit uses real-world data [3, 4, 5] to produce all results herein and can easily be modified to extend these results. Details of the toolkit can be found in Appendix A and its limitations are examined in Appendix B. Additional resources can be found online at [6].

Chapter 1

Introduction

1.1 The current spectrum crisis

In the 1920s, wireless systems were suffering heavily from excessive amounts of interference due to a lack of coordination [26]. This interference caused unnecessary performance degradation because transmitters were operating on overlapping frequencies. Congress responded by passing the Federal Radio Act of 1927 [7] (later replaced by the Federal Communications Act of 1934 [8]). These acts created the Federal Communications Commission (FCC) to oversee and facilitate spectrum allocation. They primarily use exclusive-use allocations to keep inter-system interference to a minimum [30].

These exclusive-use allocations helped to overcome the problems faced in the 1920s and they remain in place today. However, non-exclusive allocation is also an extremely important modality because it reduces barriers to entry for new devices. While large companies can afford to invest in spectrum, smaller companies and start-ups are forced to find cheaper options such as unlicensed bands. In exchange for tolerating interference from other unlicensed devices and accepting a power cap, they are allowed to freely use the spectrum. Wireless routers, cordless phones, Bluetooth devices, and remote controllers are just a few of the devices that currently operate in unlicensed spectrum. However, as we can see in the color-coded chart of Figure 1.1 that only a small portion of the spectrum is allocated for such use. In particular, these are the industrial, scientific and medical (ISM) radio bands and they amount to 276.41 MHz below 6 GHz[1]. Higher available frequencies are essentially unusable for all but short-range applications with today's technology due to their extremely poor propagation characteristics.

While these ISM bands serve their purpose well, it is common knowledge more and more wireless devices are being deployed every year. At some point, the bands will become too crowded to remain useful. There are two obvious solutions to this problem: (1) decrease usage of existing unlicensed spectrum via stricter regulations or (2) increase the amount of available spectrum. The following example illustrates the drawbacks of the first solution:

> Consider a device such as a wireless router which uses the unlicensed ISM bands and becomes so prevalent that it dominates these bands so thoroughly that other devices are unable to flourish. While it may seem reasonable for the FCC to force the devices to acquire and move to licensed spectrum in order to reduce congestion in the ISM bands, such actions would actually limit *future* use of the ISM bands. Since the device stakeholders will soon need to acquire new spectrum lest

[1] Specifically, the ISM bands are [9]: 6.78 ± 0.015 MHz, 13.560 ± 0.007 MHz, 27.12 ± 0.163 MHz, 40.68 ± 0.02 MHz, 915.0 ± 13 MHz, 2450.0 ± 50 MHz, 5.8 ± 0.075 GHz, 24.125 ± 0.125 GHz, 61.25 ± 0.250 GHz, 122.5 ± 0.500 GHz, and 245.0 ± 1 GHz. Note that $59 - 64$ GHz is also available for unlicensed used but only a small sliver (61.25 ± 0.250 GHz) is actually one of the ISM bands.

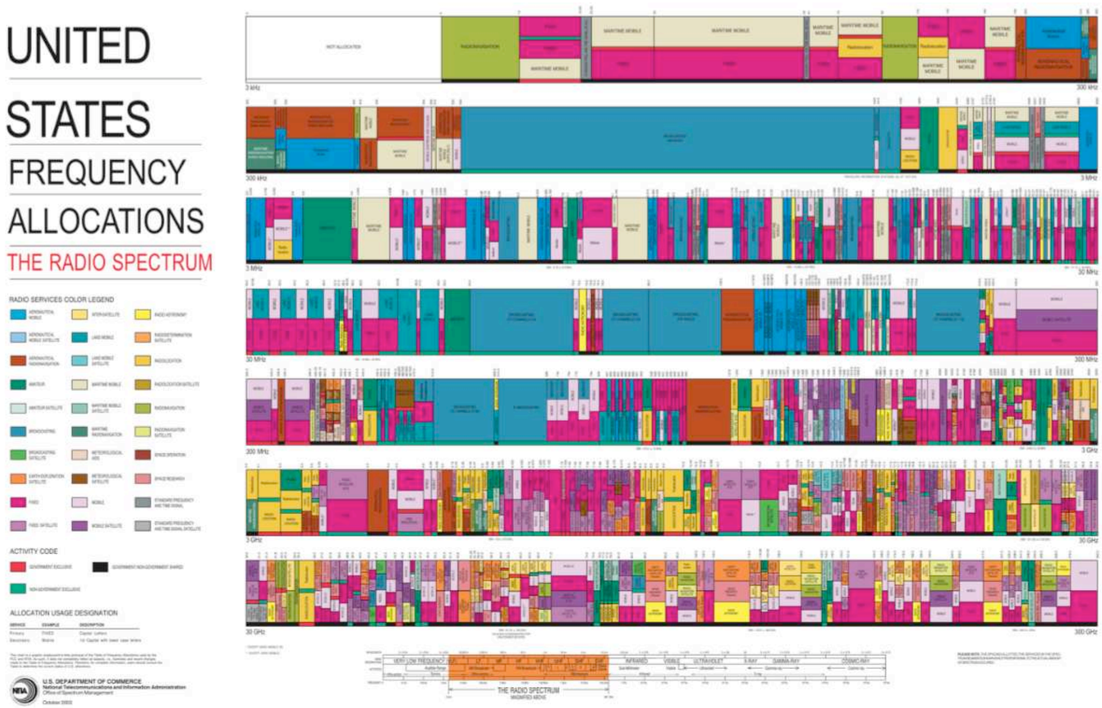

Figure 1.1: 2003 NTIA (National Telecommunications and Information Administration) chart showing the allocation of spectrum within the US [9]

> the devices face an untimely “death,” existing and potential spectrum license-holders can use this knowledge to drive up the price of spectrum arbitrarily and beyond reason, ultimately causing a lose-lose situation for the stakeholders. This example would then serve as a cautionary tale for other devices using the ISM bands: “if you want to use unlicensed bands, you cannot become “too successful.”’

Since the first “solution” could easily backfire, we need to find a way to increase the amount of available spectrum. However, Figure 1.1 suggests that all usable spectrum has been assigned. The traditional solution has been to reallocate portions of the spectrum but research suggests that there might be more clever ways of squeezing more utility out of the spectrum that is already assigned. This seems like a promising approach.

Both Taher, et al., [49] and Cabric [22] have shown that large portions of the spectrum are actually unused much of the time. Specifically Cabric showed that up to 70% of (time, location, frequency) triplets were unused. Figure 1.2 shows the measurements from this study which were taken in 2004 at the Berkeley Wireless Research Center in Berkeley, California. Orange represents occupied spectrum whereas green indicates that the spectrum is unused at that particular time and frequency. This underscores the fact that there is a large opportunity available for a frequency-agile transmitter.

The FCC took notice and in 2008 made the first ruling allowing spectrum sharing in the TV bands (updated in 2010) [1, 2]. Marcus [39] claims that the TV bands were chosen for four reasons:

- Better propagation characteristics in the TV bands than in existing Wi-Fi bands leading to both a greater range in rural areas (advancing the National Broadband Plan) as well as better building penetration.
- TV receivers require a relatively high SNR for reception which will make detection easier.

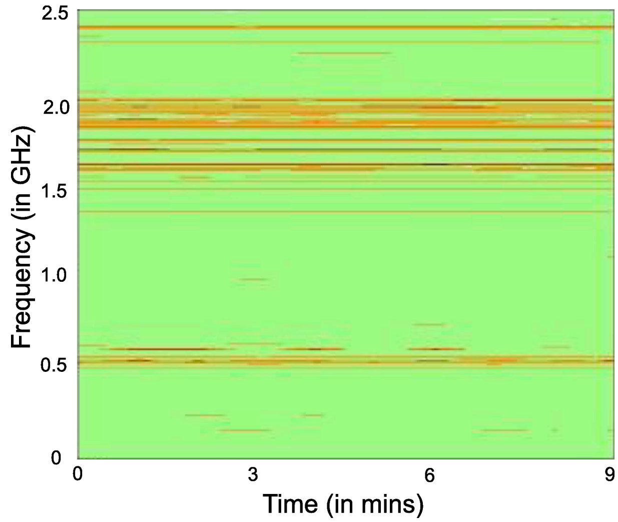

Figure 1.2: 2004 measurement of spectrum usage by the Berkeley Wireless Research Center [22]

- The 6-MHz bandwidth (per channel) is particularly attractive (as opposed to the AM or FM bands which have bandwidths of 10 and 100 kHz, respectively).
- Few people use over-the-air television so accidents have a small impact.

Finally, the FCC's vision is for spectrum sharing to spread beyond the TV bands:

> "Our actions here are expected to spur investment and innovation in applications and devices that will be used not only in the TV band but eventually in other frequency bands as well." [2, ¶1]

In order to see if the FCC's regulations will accomplish these goals, we need to evaluate the impact of their choices. To do this, we will first discuss spectrum sharing in general in the following section and then describe the FCC's choice of regulations in Section 1.3.

1.2 Spectrum sharing: an introduction

Spectrum sharing allows many heterogeneous devices under the control of different entities to operate devices on the same set of frequencies. The TV whitespaces are a special case of spectrum sharing. We will now discuss spectrum sharing in general, including the associated difficulties, and relate it to the TV whitespaces.

Typically, a set of frequencies, called a *band* or a *channel*, is licensed via auction by the FCC[2] to a corporation or agency. This entity is given exclusive rights to transmit using that frequency as well as a (implicit) guarantee (and the natural restriction that comes with this guarantee) that there will not be excessive interference from other systems using nearby frequencies.

In the spectrum sharing scenario, these entities, now called *primaries*, share their band on the basis of time, space, and/or frequency. For example, a company which has purchased nationwide rights to a band may decide to allocate some or all of this band to be used by another company (called the *secondary*) within Montana on the weekends. The arrangement may also be more flexible, as in the case of emergency bands:

[2]The following discussion is limited to spectrum sharing as it applies to the United States.

they are seldom used but at unpredictable times. In this case, secondaries may be allowed to transmit at any location within the US so long as the emergency band is not needed nearby.

It is important to note that in this scenario the two groups of users, primaries and secondaries, have fundamentally different rights. The primary has a minimum quality-of-service (e.g. bit-rate and/or locations available) that must be maintained. The secondary user is allowed to transmit within the guidelines set forth by the primary which are intended to reflect the primary's requirements.

Peha noted that spectrum sharing arrangements can be broadly grouped into two categories: cooperation and coexistence [44]. In the former, the primaries and the secondaries work together symbiotically. In the latter, secondaries endeavor to ensure that primaries are relatively unaffected by the existence of secondaries without help from the primaries. While it may seem as though the cooperative model is preferable for all parties, Peha noted that it has some drawbacks:

- Potential incompatibility with legacy equipment. If the primary user that needs protection is a "dumb" or passive device (e.g. television set), it cannot communicate to the secondary to make it aware of misbehavior even if it wanted to do so.
- Additional constraints on secondary devices. It is desirable to provide maximum flexibility to both parties. Forcing cooperation between primaries and secondaries places additional constraints (e.g. that they must be able to intercommunicate – or communicate at all, since some secondaries may prefer to be passive devices) on these devices which may not be justified by the increase in quality of service.
- Interaction may consume too many resources. Even if well-engineered, interaction between primaries and secondaries may require too much time, power, or other overhead to be justifiable.

In 2008, the FCC released rules for spectrum sharing in the television bands under the coexistence[3] model [1] (updated in 2010 [2]). Specifically, these regulations allow for unlicensed devices conforming to the FCC's restrictions (which include power, height, and location) to transmit using the over-the-air TV broadcast bands. These restrictions are discussed in further detail in the following section.

1.3 Overview of current regulations

In this section we will briefly review the relevant portions of the regulations set forth by the FCC in Appendix B of [2] for TV whitespace devices which utilize databases to find "unused" spectrum. The rules also permit the operation of sensing-only devices, as discussed in Section 1.4.2.1.

Secondary devices are classified as depicted in Figure 1.3a. Restrictions on a device are based on its classification, as shown in Figure 1.3b. Restrictions generally fall into a few categories:

- **Power**: a limit on the effective isotropic radiated power (EIRP) or on amount of power delivered to the antenna and the directional gain.
- **Height**: a limit on the height above average terrain (HAAT) of a device. This is only imposed for fixed devices.
- **Channels**: no secondary device will be allowed to transmit in channels 3, 4, and 37. Channels 3 and 4 are used by legacy equipment such as VCRs and early DVD players. Channel 37 is not used for TV broadcasting; instead, it is reserved for radio astronomy and wireless medical telemetry services. This

[3] Although the rules are predominantly in the coexistence category, they do require some cooperation by proxy: TV stations and wireless microphones are registered *by their owners* with the TV whitespace databases.

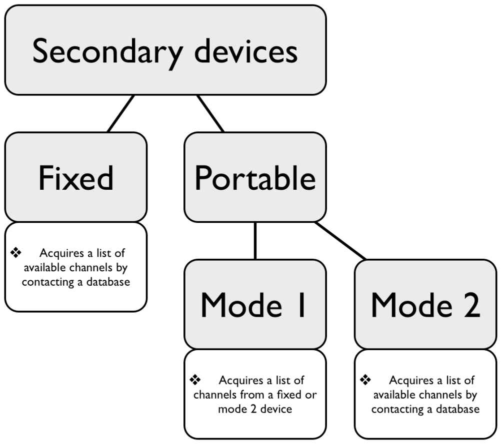

(a) Classification

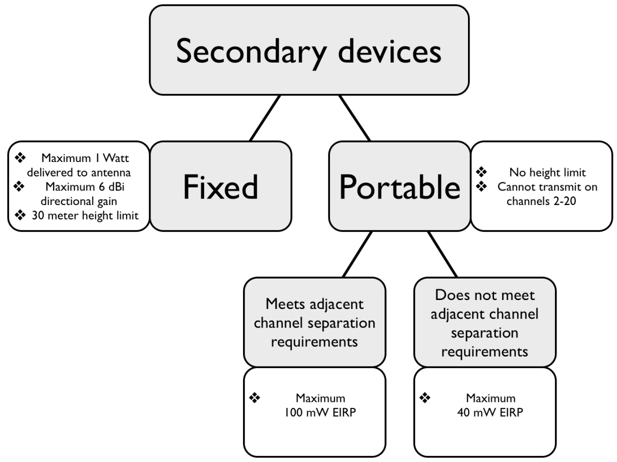

(b) Restrictions

Figure 1.3: Device classification and restrictions for database-connected devices according to the 2010 FCC rules [2].

leaves the remainder of channels 2-51 for secondary use[4]. Further restrictions on channel usage are imposed based on location, as described below.

- **Location (cochannel)**: no secondary devices of any kind are allowed on the same channel within the protected radius (called r_p) of a TV tower. As the adjective "protected" implies, the TV signal should still be decodable within this region. This radius is a function of the characteristics of the TV tower. Furthermore, each device must maintain an additional separation distance, $r_n - r_p$, from the edge of the protected region. This value ranges from 6.0 km to 14.4 km depending on the height of the secondary device. This is shown in Figure 1.4 as the "no man's land."

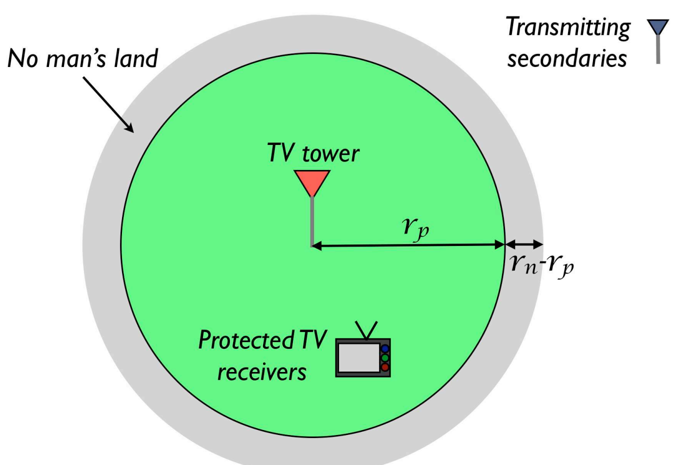

Figure 1.4: Illustration of the protection radius r_p and the no-talk radius r_n.

[4]The FCC also bans transmission on channels 2-21 for portable devices.
Furthermore, it reserves two channels at all locations for wireless microphone usage. For example, channels 24 and 41 are reserved near Chicago and channels 15 and 16 are reserved in Berkeley.

- **Location (adjacent channel)**: secondaries are also excluded from the circle of radius $r_{n,A}$ centered at the TV transmitter in the adjacent channel. The protected radius is the same as in the cochannel case; however the value of $r_{n,A} - r_p$ again depends on the height of the secondary transmitter and varies from 0.1 km to 0.74 km. Note that $r_n - r_p \neq r_{n,A} - r_p$, so $r_n \neq r_{n,A}$. A secondary which is outside of every tower's $r_{n,A}$ on all adjacent channels is said to meet adjacent channel separation requirements.

For simplicity, we consider only fixed devices in our model but our results are generalizable to the portable case[5]. We also assume that each secondary has perfect knowledge of his location and the channels available for his use.

We also do not consider the specifically excluded regions such as radio astronomy protection zones, but these are generally areas of low population and thus have little impact on our study. The current FCC database only lists towers in the United States, thus we also ignore TV transmitters in Canada and Mexico[6].

Further limitations of our model are discussed in Appendix B.

1.4 Related work

Research in cognitive radio can be roughly divided into two components: technology and policy. Policy research investigates the design and enforcement of regulations while technology research explores the design and operation of the devices. Of course, technology research informs policy research and vice versa. Some of the major branches of work are shown in Figure 1.5 and many of them are discussed in the rest of this section. Our work concentrates on defining regulations and analyzing them, shown by the bold blocks in Figure 1.5.

On the policy side, we will discuss the challenges associated with whitespace regulation and in particular the importance of well-defined property rights. Finally, we will look at the various ways that people have attempted to quantify the utility resulting from current defined and proposed regulations. This is an important component in the feedback system: through these simulations we can understand how to improve regulations.

With regards to technology, we will concentrate on the two different ways of finding whitespaces: sensing and databases. We will highlight the strengths and weaknesses of each and give examples of how to overcome some of their associated challenges.

1.4.1 Policy

Policy research involves a feedback loop between defining regulations and evaluating them. In this section we look at current and proposed regulations as well as some attempts to quantify the quality of these regulations.

1.4.1.1 Defining the rules

Current rules

[5] A change in propagation model is required since the ITU model [10] does not accommodate transmitters with heights less than 10 meters.

[6] The FCC rules state that the protected contours of non-US TV towers must be protected within the country of origin up to the US border, but not beyond. Thus US residents who previously received Canadian or Mexican TV stations are not guaranteed to continue to receive these stations.

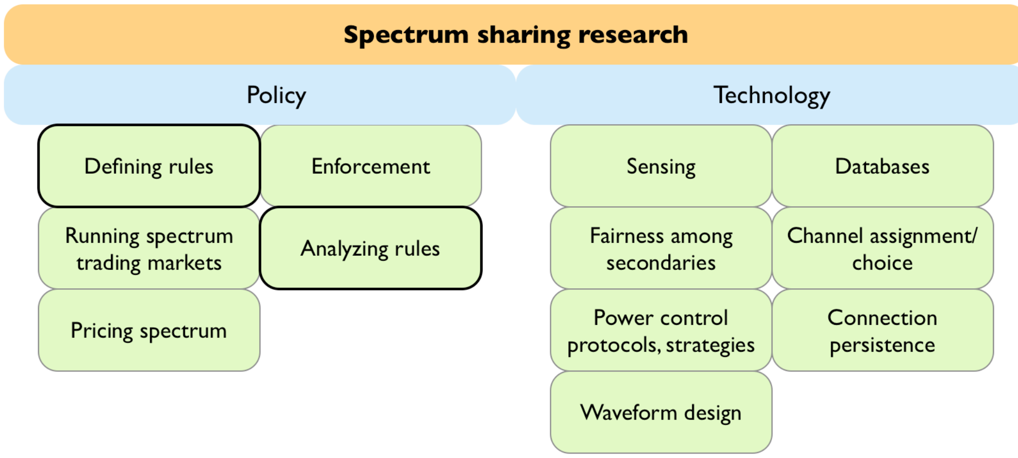

Figure 1.5: Current topics of research in cognitive radio

As mentioned previously, in 2008 the Federal Communications Commission (FCC) defined the first rules for the use of TV bands devices (TVBDs, i.e. cognitive radios in the TV bands) within the United States [1]. In 2010, they updated these rules to reflect decisions made in response to petitions from a variety of interested parties, including Adaptrum, Motorola, Dell, Microsoft, and Google [2]. Briefly, these rules allow devices to transmit on the same channels as TV stations but with uniform limitations on their emitted power, height, and choice of channel. A more detailed summary can be found in Section 1.3.

Regulatory bodies around the world are in the process of adopting or improving similar rules: Ofcom in the United Kingdom [11], the Electronic Communications Committee of Europe [19], and the Australian Communications and Media Authority [12].

Proposed rule modifications

There have been many suggestions of alternative rules. In this section, we present one which addresses interference aggregation as it relates to power scaling for secondaries. Interference aggregation occurs when signals from secondaries add up at the TV receiver. Approaches such as those adopted by the FCC consider a single secondary transmitter and thus neglect this effect. As we will show in Chapter 4, it can be a big mistake to neglect the effects of interference aggregation.

The work discussed below accounts for interference aggregation while suggesting a potential power scaling rule which would allow secondaries to use a different maximum power based on their time, location, and operating frequency. This has the potential to increase secondary utility while maintaining the desired quality of service for the primary.

To get a handle on how a power scaling rule will play out and to set up their model, the authors first define a secondary interference constraint based on system parameters such that primary receivers inside of r_p will be protected. They then determine the maximum "safe" (i.e. meets the interference constraint) power for a single secondary located distance $r_p + d$ from the TV transmitter as a function of d. This "safe" power scales quickly with d and in fact "jumps from zero to health-endangering levels as soon as the secondary transmitter is outside the protected radius" [36]. The FCC appears to have used a similar approach when designing their rules, ultimately choosing $d = r_n - r_p = 14.4$ km[7].

But the whitespaces were not created for a single device and their rules should not be designed with that assumption. Recognizing this, the authors also derive upper and lower bounds on a "safe" *power density*[8] using strict subsets and supersets, respectively, of the potential secondary transmitters.

[7]Actually, the value for d varies depending on the height of the transmitter and ranges from 6.0 to 14.4 km.

[8]We envision each secondary transmitter using power $P = D \cdot A$ when the local power density is D and the transmitter has footprint A. This notion is developed further in Chapter 4.

If we assume a standard theoretical inverse-power-law path-loss model ($P_{\text{received}} = P_{\text{transmitted}} \cdot (\text{distance})^{-\alpha}$ where $\alpha > 2$) , we additionally see that the multiple secondaries can in fact be viewed as a single transmitter located at $r_n - r_p$ with new path-loss exponent $\alpha' = \alpha - 2$. This is surprising because it indicates that the interference term is not dominated by the transmitter nearest the TV receiver[9].

Finally, Hoven shows in [35] that decreasing the value of $r_n - r_p$ will have a significant effect on secondaries located far from r_p: allowing transmissions nearer to the primary receiver means that the power density everywhere needs to be decreased. In essence, secondaries far from r_p are sacrificing some utility for those near r_p but *outside* r_n. At the same time, secondaries *inside* r_n are sacrificed (since they may not transmit on this channel due to our choice of r_n) for the sake the far-away secondaries. It is not immediately clear how to choose r_n, though Vu, et al., have done some preliminary work on this subject [55]. This tension will be made more explicit and subsequently partially mitigated in Section 4.

This work was also done for generic path-loss models. Furthermore, the authors only examined one power scaling rule but point out that many valid power scaling rules do exist. We will further discuss the problem of interference aggregation on a national scale and develop similar power scaling rules in Section 4.

Fairness in the rules

The FCC's goal in creating the whitespaces is explicit:

> "Our actions here are expected to spur investment and innovation in applications and devices that will be used not only in the TV band but eventually in other frequency bands as well." [2, ¶1]

Furthermore, they recognize that in order to attain these goals they must allow the maximum amount of flexibility in the requirements for whitespace devices. For example, when imposing limits on secondary power-spectral density (PSD), they state

> "We decline, however, to adopt minimum bandwidth requirements as requested by IEEE 802 and SBE. We find that a minimum bandwidth requirement could unnecessarily constrain the types of modulation that could be used with TV bands devices and is not necessary because the PSD limit has the same effect of preventing high power levels in a TV channel." [2, ¶83]

Despite their laudable efforts to maintain flexibility in secondary requirements, they have fallen victim to the sensing-only paradigm and consequently overlooked an important detail: now that we have introduced databases, a tremendous number of new possibilities have emerged. The sensing-only paradigm is outdated and overly restrictive: devices are already required to contact a database with high frequency and there's no reason a power limit can't be included in the information obtained from the database[10]. In fact, [38] has already suggested this modification.

As pointed out in the previous section, Hoven and Sahai's work [35, 36] shows the FCC's rules are a degenerate case of a power-scaling rule. That is to say, they could have chosen a more flexible power rule but instead picked the tried-and-true uniform power limit.

Most importantly, this uniform power limit requirement carries with it the implicit choice of application. For example, under the current rules the whitespaces are unappealing to long-range applications due to their low power limit. However, increasing the power limit requires increasing the size of the "no-man's-land" (also

[9] Woyach, Grover, and Sahai discuss this in [56]. They show that the aggregate interference behaves qualitatively differently depending on the range of the transmitters to the offended receiver. In particular, they show that the distribution of the aggregate interference changes from approximately Gaussian to a heavy-tailed distribution when randomly-placed interferers approach r_p. Regulators may need to incorporate this difference into their regulations, depending on the situation.

[10] A form of power scaling is already required for secondaries: *"TVBDs [i.e.. secondaries] shall incorporate transmit power control to limit their operating power to the minimum necessary for successful communication. Applicants for equipment certification shall include a description of a device's transmit power control feature mechanism."* [2, §15.709(a)(3)]

called d and $r_n - r_p$) which may reasonably be opposed by devices for which the power limit is sufficient. This tension between types of users will be further explored in Section 4. However, this can be viewed as one of a family of fairness problems. This tension is present in all systems and it's important to understand how others have solved this problem in smaller, simpler systems.

There is an extensive body of work on fairness in networks [51, 20, 45] but relatively little work directly discussing fairness in the whitespaces. Many economists such as Coase believe that the optimal strategy is to define property rights and a means by which they may be exchanged (i.e. a trading market) and let stakeholders trade rights until they are happy [24]. In the next section we will discuss the conditions necessary for this strategy to lead to the optimal solution.

1.4.1.2 Spectrum property rights and trading markets

Property rights define the guarantees and restrictions that come with owning or leasing that property. Strong definitions are important in preventing inefficiencies, both at run-time and in the courtroom. We will first provide examples which uphold this claim and then give an example of a set of property rights which would have prevented the problems showcased in [27]. Finally, we briefly discuss some property rights through the lens of enforcement.

Motivation

An excellent example of the importance of well-defined property rights is given in [26] and is repeated below.

> *Example:* Two oil companies, A and B, are pumping from the same reservoir. Since A's use decreases B's opportunity and vice versa, each will pump as quickly as possible. In order to do so, they will end up consuming the resource inefficiently: to rid themselves of the excess, they will flood the market by selling at a lower rate.
>
> *Solution 1:* Assign exclusive rights to either company A or B. Now the owner has incentive to use the resource more efficiently. This is essentially what the FCC has done for many years. However, companies may still choose to waste a resource (e.g. TV bands in Montana) because it is not profitable enough for them. The next solution fixes this problem.
>
> *Solution 2:* Establish quotas on how much of the resource each person can use. In the cognitive radio setting, this may mean setting limits on any of the following: time, location, frequency, transmitter height, and interference level.

In fact, De Vany, et al., even point out that

> "It was the nonexclusivity of spectrum-use rights during 1925-1927 that initially produced government regulation by the Federal Radio Commission and later the Federal Communications Commission." [26]

Finally, [27] examines three real cases in which poorly-defined spectrum rights led to very costly outcomes for taxpayers and companies alike.

Example 1: both parties obeyed the rules but still conflicted (800 MHz)

- Problem: public safety receivers were unable to operate effectively with Nextel in a nearby band.
- Resolution: Nextel "paid" the costs to reconfigure the 800 MHz band in exchange for 10 MHz of contiguous spectrum so fault was never determined; because of terms of the settlement, the US government essentially ended up paying for the reconfiguration (up to $2.8B).

- Why it happened: "There had not been a sufficiently clear delineation of who had what rights according to their individual licenses, and what actions each party could be required to take to resolve the conflict" [27].
- Similar case: T-Mobile unexpectedly caused interference in a nearby band; their obligation was poorly defined and caused a two-year delay while they installed equipment for the protected party.

Example 2: FCC changed rights in the middle of the license period

- SDARS (e.g. Sirius XM satellite radio) was allowed to install repeaters under an *experimental* license.
- WCS (in an adjacent band with no guard band) was worried about interference from SDARS devices and successfully lobbied the FCC to relax out-of-band emission limits for their band.
- Sirius XM argued against the changes for WCS because it affected their networks in which they had heavily invested despite having declared their system an *experiment*.
- Moral: the FCC can and will change the rules at any time. This case serves as an important example to stakeholders which commonly assume that the rules are fixed.

Example 3: lack of clarity led to protracted discussions

- M2Z offered to build a free nationwide broadband service in exchange for 20 MHz of spectrum with a 15-year license.
- The FCC declined their application, initiated a rulemaking, and asked for public comment.
- T-Mobile objected saying that M2Z's TDD systems would interfere with their own FDD systems.
- After almost four years of comments, the FCC closed the rulemaking and finally informed M2Z of its decision.

These examples show three ways in which poorly-defined rights can be costly, both monetarily and temporally. The next section suggests a set of property rights which would have prevented some of these problems.

Defining property rights

De Vany, et al. [26] have published an extensive report regarding the property system required for spectrum sharing. Specifically, they suggest that the property system must contain three elements:

1. Precise definitions of rights which are both unambiguous but also compatible with the properties of electromagnetic radiation
2. Mechanisms for legal enforcement of rights
3. Methods for transitioning the spectrum rights

Just as in the case of water rights, we cannot isolate the actions of any given user. Radio waves propagate far beyond their point of origin and thus affect a potentially large number of parties. We must determine how to define the rights of all parties rigidly enough that they can be easily checked and enforced but flexibly enough so that the resulting property will be useful.

The authors suggest a particular set of property rights, called a *TAS package*. The TAS package would define both the time a party was allowed to access the channel and which channel (frequency) he was allowed to use. In particular:

- The owner of a TAS package has two area rights.
    1. The exclusive right to transmit so long as his signal strength does not exceed the given threshold outside of his area
    2. The right to be free of interference exceeding this same threshold within his area
- The owner of a TAS package has two frequency rights.
    1. The exclusive right to transmit so long as his signal strength does not exceed the given threshold outside of his frequency band
    2. The right to be free of interference exceeding this same threshold within his frequency band

De Vany, et al., also argue that the area referenced above should be defined with straight lines rather than circles[11] to reduce exchange costs: such definitions would make subdivision straightforward. Additionally, non-overlapping circular properties cannot cover the entire area; if "squatters" begin using the unassigned regions, enforcement may become more complicated.

While it may seem as though property rights are simply a pedantic exercise, they can actually have far-reaching implications for the systems which use them. For example, Gastpar examined achievable capacities with received-signal-strength constraints similar to those given above. He showed that neither feedback nor collaboration can increase the capacity region of an additive white Gaussian MAC under this constraint [29]. This result is surprising because both feedback and collaboration increase the capacity region under the typical per-device maximum-power constraint [52].

Despite this undesirable result, such property rights are quite useful. The following section discusses the problems inherent in other popular property-right proposals.

Enforcement policies

We must also provide mechanisms to enforce these rights easily and fairly; if done incorrectly, a great many costly and difficult-to-decide legal battles will be fought (examples above courtesy of [27]). For example, consider the following three schemes (rejected by [27]):

1. Rule: "don't cause harmful interference." While it is easy to agree with the sentiment of this rule, it unfortunately does not specify receiver characteristics (e.g. band rejection) so one can always find a receiver such that "harmful interference" is caused. Specifying these characteristics is contrary to the light-handed regulation style that makes property rights both valuable and easy to trade.

2. Rule: new licenses are planned so that they do not degrade neighboring services. This is an understandably feeble guarantee in the eyes of the incumbent since the FCC has already demonstrated its propensity to alter licenses in the middle of their term, not even just at auction time.

3. Rule: receivers are guaranteed that interference experienced will not exceed a certain threshold (similar to De Vany, et al., in [26]). However, it can be difficult to determine fault in the event that the threshold is exceeded.

De Vries suggests fully removing the idea of "harmful interference" since it is too vague a term, just as De Vany, et al., did in [26]. Instead, they suggest declaring a party guilty if they exceed their emissions requirements regardless of intent. The beauty of such a system is that enforcement can be done automatically without the need for protracted arguments or settlements. De Vries further asserts that this will aid regulators in determining whether or not a new set of rights can coexist with existing rights.

Other issues such as defining a way to exchange these property rights as well as enforcing them "in the wild" [58, 57] are also important. In this thesis we ignore the challenge of enforcement by assuming that the FCC will only certify those devices which comply with database instructions.

[11] The contour lines for the signal strength from an omnidirectional antenna in an empty landscape are circular, thus circular properties would more closely match this natural phenomenon. The authors suggest the use of directional antennas to minimize wasted potential.

1.4.1.3 Analyzing regulations

It is important to analyze potential and existing regulations in order to verify that they will function as intended: the examples in [27] (also described above) show that poorly-defined regulations can have large monetary and temporal costs. Some of these costs may even go unnoticed in the form of services which are never even created because there is no viable spectrum option. Additionally, analysis of such regulations may uncover improvements which advance the original goals of the regulation.

There are two ways the analyze the rules: (1) using simple examples and single-location measurements to try to draw conclusions about the big picture; and (2) applying these simple examples to the big pictures. We will discuss each in turn.

Initial studies

Taher, et al., collected multiple years of spectrum observations in Chicago, Illinois [49]. They first present their results in the form of a novel color-coded image which highlights spectrum usage trends (e.g. digital transition, holidays, daily and weekly peaks, etc.). The authors also provide aggregate statistics, perhaps the most shocking of which is the 14% spectrum utilization in the 30-3000 MHz frequency range for 2010 (through October).

The New America Foundation provides an estimate of the number of channels available for secondary use in major metropolitan areas [23]. Their estimates appear to be optimistic (for example, they predict the availability of 19 channels in San Francisco) but they do not detail their methods so we cannot explain the difference between their results and our own (Section 2.1).

Comprehensive studies

The authors of [43] were the first to examine the white space opportunity in the United Kingdom (2009). They used TV coverage maps to predict the protected radius (r_p) but they do not include the no-man's-land ($r_n - r_p$). When ignoring adjacent-channel exclusions, they found that an average of approximately 150 MHz are available 18 major population centers in the UK to low-power transmitters. This number drops to about 30 MHz when including adjacent-channel exclusions.

Mishra, et al., were the first to do a detailed study of the whitespace opportunity in the United States [41, 40]. They used real-world data to determine how many channels were available to secondaries in 2009 under the FCC's 2008 rules. They then extended the question by asking how many channels are available if a secondary has a minimum SNR requirement (TV signals propagate far beyond r_p and effectively raise the noise floor for secondaries). Finally, they looked at the impact that the FCC's choice of rules has on the number of channels available. Notably, they have also posted their code online at [13]. The work in this thesis builds upon this body of work[12].

Van der Beek, et al. performed a similar experiment for Europe [53], finding that the average person is allowed secondary access on 49% of TV channels in Europe. This confirms their suspicion that whitespaces are more plentiful in the United States than in Europe. They then analyze the sensitivity of these results to the use of terrain models and various propagation models. They find that while the aggregate statistics remain roughly the same, changing the propagation model tends to affect local results.

Jäntti, et al., conducted a similar study for Finland, finding that the number of available channels varies from 4-5 in the south to 35-40 in the north which is less populated [37]. They also evaluate achievable data rates using a simple model and claim that the ECC's proposed rules allow for higher data rates than those possible under the FCC's rules.

Finally, Microsoft's WhiteSpaceFinder [14] and Spectrum Bridge's "Show My White Space" [15] offer everyone the chance to see the amount of whitespace available at a given address. However, this design choice obscures the big picture by focusing on single addresses rather than showing results for an entire area.

[12]A point of history: original results by the author of this thesis, published in [33, 34, 47], extended the code provided in [13] but used the same external data. This thesis includes updated population and tower data, as discussed in Appendix A.2. Both use the same propagation model [10].

1.4.2 Technology

Technology research for the whitespaces deals primarily with technical issues regarding device operation. Two such challenges are discussed below: channel availability estimation and reliable power control for the sake of primary safety.

1.4.2.1 Channel availability estimation: databases versus sensing

Under the FCC rules, a cognitive radio is afforded two methods by which he may determine which channels are available to him: sensing and contacting a database (directly or indirectly). This section first describes the rules for each method and then discusses the technological challenges associated with each.

Sensing

Secondaries using the traditional sensing approach must be able to detect TV signals with a sensitivity of -114 dBm. This sensitivity helps guard against scenarios such as the hidden terminal problem in which a primary transmitter-receiver pair and a secondary transmitter are arranged in a triangle with a bad fade between the primary transmitter and the secondary transmitter. Thus, the secondary transmitter may not notice that there is a TV transmitter there, think it's OK to transmit and ultimately ends up interfering with the primary receiver.

In [41, 40], the authors use real TV tower data for the United States to show that sensing can drastically decrease the opportunity available to secondaries as compared to the database approach. Specifically, the average secondary user with database capabilities may use 22.43 whitespace channels whereas the average secondary user that only has sensing capabilities and adheres to the FCC's -114 dBm sensing requirement may use only 9.8 channels.

In a follow-up paper the authors showed that cooperative sensing can help overcome this problem [50]. Specifically, diversity provides much of the gain experienced over individual sensing. However, the potential for correlation of samples across users means that regulators still need to be conservative in their choices.

Goncalves, et al., expand on this idea in [31], arguing that distributed sensing can achieve similar results to the database approach with similar monetary and energy costs. As they point out, there is a tendency in the community to think only about the database approach and disregard sensing as a bad idea which requires too much overhead.

Databases

Databases represent a paradigm shift: with the near-ubiquity of Internet access in the modern day, we no longer need to rely purely on sensing. Whitespace devices can instead query a database in real-time via the Internet in order to determine the set of allowed channels. Advantages to the database approach include more efficient use of spectrum [32] as well as the potential for power scaling which further increases the utility of the available spectrum [38]. These databases present an enormous opportunity which is currently in its infancy and whose limits must be more thoroughly explored in the coming years.

Secondaries who wish to connect to the database must be able to pinpoint their location to within 50 meters[13]. Once determined, they transmit their location (method unspecified) to a TV whitespaces database which will tell them the number of channels locally available. Ten database administrators were selected by the FCC in January of 2011, including Google, Neustar, and Spectrum Bridge [16].

[13] Secondaries without the ability to directly contact the database are allowed to contact a secondary which has a direct connection to a database. The latter will contact the database on behalf of the former and return the results.

One of the drawbacks to this approach is that location uncertainty is multiplied. If the directly-connected secondary is near the edge of r_n, we may end up allowing the indirectly-connected secondary to transmit at a distance well inside of r_n and perhaps even inside of r_p. The FCC rules are a bit unclear about just how far away the indirectly-connected secondary may be from the directly-connected secondary in order to request a channel list. Here, the SNR for the connection between the two could act as a proxy for distance, though even that is prone to errors given that the noise levels will be quickly varying precisely where we anticipate a problem.

Although the drawbacks to databases have not been extensively studied, a few are readily apparent. The most obvious is that there are only a few points of failure (whether they are hacked or simple offline) which can lead to catastrophic outcomes. They also rely heavily on propagation models. To that end, Microsoft researchers showed the importance of a good propagation model and accurate knowledge of secondary location [42].

1.4.2.2 Power control protocols, strategies

The potential rules discussed earlier and those that we will discuss in Chapter 4 require secondaries to be able to dynamically scale their power. This is an old problem in traditional cellular systems [30]but this research does not immediately apply to the problem of power control in the TV whitespaces. This in part due to the strict safety requirements of the secondary. Two approaches are detailed below.

Pollin, et al., suggest a distributed power control algorithm which allows secondaries to safely scale their power by estimating their distance to the protected receivers and the secondary's contour (i.e. region of interference) [46]. Vanwinckelen, et al., extend this work to show that it provides a reliable method for power control [54]. However, this work neglects the effects of aggregate interference.

Bater, et al., present a centralized method of allocating power across transmitters while respecting the primary's interference constraint [21]. These receiver-centric constraints actually lead to increased spectrum utilization but they are expected to have higher overhead than transmitter-centric methods due to the increased number of involved parties.

1.5 Contributions

This thesis builds conceptually on the work done by Mubaraq Mishra in [40, 41, 13]. It also extends his work [13] in building a toolkit for studying the TV whitespaces which incorporates real-world population data [4, 5], TV assignment data [3], and an empirical propagation model[14] [10]. This toolkit offers a new way of evaluating the TV white spaces as well as potential changes to regulations. The toolkit and its documentation are freely available online at [6].

Using this toolkit, we have not only quantified the amount of raw spectrum available but we have also shown its potential utility (achievable data rate) via models that incorporate varying ranges of secondaries, economic realities, and the effects of both noise from primaries and self-interference. Studies [43, 40, 41] prior to our initial publication on this topic [33] did not consider this metric but subsequent studies [37, 53] did. Furthermore, the Australian Communications and Media Authority is advising that a similar study be conducted for the whitespaces in Australia [28].

We then compare two simple changes to the rules: altering the existing TV protections and straightforward channel reassignment. We show that whitespaces are the only rational option if the incumbent service is to be preserved to any reasonable degree. Mishra, et al., performed similar work but again used only the number of channels as a metric for secondary utility [40].

The FCC's regulations were unrealistically designed with a single secondary in mind. We show that the aggregate interference coming from secondaries in a realistic deployment may be sufficient to impact TV reception. We suggest introducing a maximum power density rule – now possible with the advent of database-controlled devices – to solve this problem.

Finally, we suggest a principled way by which a maximum power density can be set in order to compromise between rural and urban users, a feature not found in the current FCC regulations and clearly desired by entities such as PISC (Public Interest Spectrum Coalition), WISPA (Wireless Internet Service Providers Association), and Motorola [2, ¶71] . We explore the implications of such a rule and show that indeed secondaries fare well under our suggested power density rule.

[14] Note that both the population data and the TV assignment data have been updated from their use in [40, 41, 13]. The propagation model remains the same.

Chapter 2

Evaluating the whitespaces

This section quantifies the opportunity available to secondaries obeying the FCC's regulations in the TV whitespaces for both single-link and cellular applications. It also determines the relative impact of various effects: the no-talk regions necessary for primary operation, noise from primaries, self-interference from secondaries, height, and range.

The FCC created exciting new prospects when it designed the TV whitespaces, the "vacant" areas of the TV bands in which unlicensed devices are permitted to operate. These devices are allowed to transmit when they are far enough from the TV towers. We first explore the amount of spectrum the FCC regulations provide for unlicensed use. However, this spectrum is not actually "white" due to pollution (TV signals which are essentially noise to secondaries), making it less valuable than clean spectrum. We will show using a simple single-link model that this pollution is actually a minor effect in comparison to the need to halt transmissions near the TV towers.

Along with environmental concerns such as TV signal noise, devices operating in the TV whitespaces must concern themselves with traditional system parameters, namely transmitter height and communication range. We will show that in general the range plays a significantly larger role than transmitter height. However, a poor choice of height may lead to a 50% degradation in link quality.

In addition to pollution, secondaries must also concern themselves with noise from other secondary devices called self-interference. We incorporate medium access control (MAC), a common method for mitigating self-interference [30]. In this scheme, devices take turns transmitting in order to increase their individual rate. Using this model we show that the single-link rates are overly optimistic.

Furthermore, although some systems operate fixed-range links, many do not. We argue that the range is often a function of the local population density rather than an endogenous variable. Under this model, we see that rural areas experience much lower rates despite having more available channels due to large communication ranges. Conversely, urban areas do quite well despite having few channels.

Finally, we consider using a standard cellular model in addition to our point-to-point model. We show that, similar to the results for the MAC model, the opportunity has decreased relative to the single-link model.

2.1 Channels available to secondaries

Since TV towers tend to be located around population center and the presence of a TV station precludes secondary use in the nearby area, it may be reasonable to conclude that there are fewer channels available in higher population areas.

In order to verify this claim, we use real tower data for the United States [3] to evaluate the FCC rules [2] on a nationwide scale[1]. We enforce cochannel and adjacent channel exclusions, as defined and discussed in Section 1.3, and then count the number of available channels at each location. This is shown in Figure 2.1a in the form of a color-coded map. We see that up to 45 channels are available in some locations and even Los Angeles has at least four channels available in most places.

Each TV channel is 6 MHz wide. For comparison, wireless routers using the IEEE 802.11 standard and operating in the 2.4 GHz ISM band use channels which are 22 MHz wide, so a wireless router using the TV white spaces would require four TV channels for operation. However, signals attenuate faster in the ISM bands than they do in the TV bands, so 22 MHz in the TV bands is much more useful than the same amount of spectrum in the ISM bands.

Figure 2.1b shows the complementary cumulative distribution function[2] (CCDF) by population for the map of Figure 2.1a. From this we can see that at least 87% of the population has at least four channels and could therefore operate a wireless router in the TV whitespaces. Furthermore, only about 1.5% of people have zero channels available (colored black on the map).

[1] Details on map generation can be found in Appendix A.
[2] Details on CCDF calculations can be found in Appendix A.

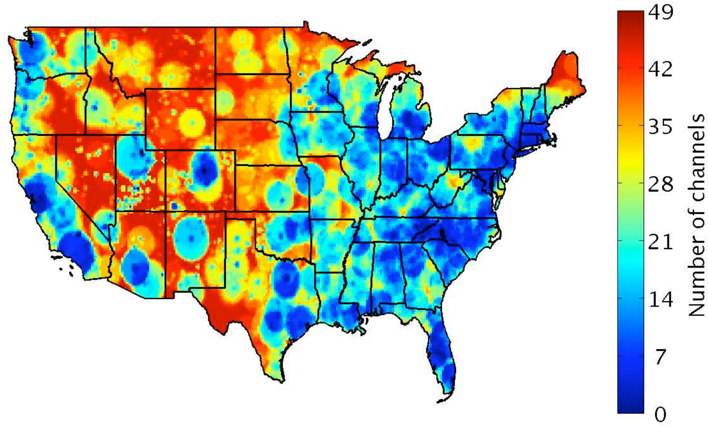

(a) Color-coded map of the continental United States

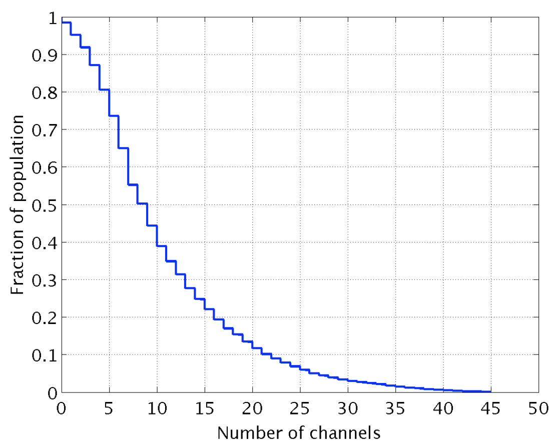

(b) CCDF by population of the number of channels available to secondaries nationwide

Figure 2.1: Number of 6 MHz channels available to secondaries in the United States

Notice that the coasts of the United States have fewer available channels than the central portion of the nation. In fact, one of the key obstacles to secondaries is that high-population areas have the least bandwidth. This phenomenon is due to the fact that TV towers (which "take" channels from secondaries) tend to be located near population centers. As we suspected, a larger population generally means fewer channels available for secondary use. We see this more clearly in Figure 2.2: at higher population densities, there are typically fewer channels available.

However, the data presented are potentially misleading. Not all whitespace channels are alike because the TV signals, which do not simply end at r_p, raise the noise floor for secondaries in a spatially- and frequency-varying manner. The next section discusses this problem and quantifies its impact.

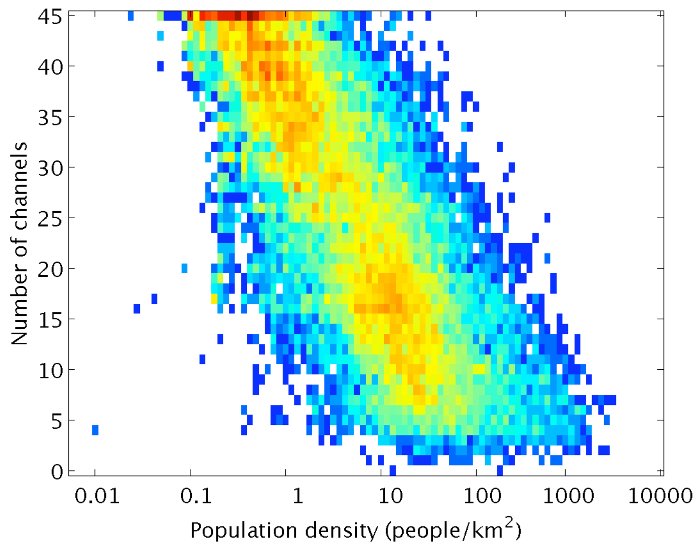

Figure 2.2: Color-coded histogram (of pixels) showing the inverse relationship between population density and the number of channels available for secondary use. Red indicates the highest frequency and blue the lowest frequency. White indicates that there were no data points in that bin.

2.2 Impact of primaries on secondary utility: are the whitespaces "white"?

A simple count of the number of channels available to secondaries is not sufficient to evaluate the TV whitespaces for secondary use. Spectrum users are interested in quantities such as deliverable data rate rather than the amount of spectrum available. In exclusively-allocated channels these quantities are linearly related; however, shared bands are quite different in this regard.

Primaries affect the achievable rates of secondary users through two factors: noise from TV towers and location-based channel restrictions. We will see in this section that the effect of the noise is small relative to that of the exclusions.

To assess the available data rate, we use the information-theoretic Shannon capacity formula [48]:

$$\text{rate} = \text{bandwidth} \cdot \log_2 \left(1 + \frac{\text{received signal (desired) power}}{\text{noise (undesired) received power}}\right)$$

The the fraction of received signal (desired) power to noise (undesired) received power is commonly referred to as the signal to noise ratio (SNR). For simplicity we do not consider schemes such as dirty paper coding [25] which use information about the noise to increase the rate.

In the TV whitespaces, the bandwidth of each channel is 6 MHz and rate can be summed across channels. Thus we can express the total rate at a given location as:

$$\text{total rate} = \sum_{c \in \text{TV channels}} a_c \cdot (6 \cdot 10^6) \cdot \log_2 (1 + \text{SNR}_c)$$

where $a_c = 1$ if channel c is available for secondary use at that location and $a_c = 0$ otherwise.

Unlike in many other bands, TV whitespace devices may be subject to substantial noise due to TV tower transmissions. We now examine the effect of this noise as well as protected areas on the achievable data rates for whitespace devices.

In typical wireless systems where the band is licensed, the license-holder is free from in-band interference (a.k.a. noise) and has a guaranteed limit on the amount of out-of-band interference to which the system is subjected. Thus, the majority of the interference that he faces is due to transmitters he has installed himself. With this level of control over the interference, he can carefully choose the locations of his transmitters to maximize his utility.

In unlicensed bands, devices operate at much lower powers. In this way, even though individual devices cannot control the amount of noise experienced due to other transmitters, they are reasonably assured that the total interference will be small.

The situation for the secondaries in a spectrum-sharing situation is essentially the worst of both worlds. They lack control over the other secondary transmitters in the band and they must cope with the signals from the primary transmitters which is essentially noise to them. Furthermore, they have to turn off whenever they get too close to the primary.

To understand how these effects play out, we look at a toy example in which we have a single secondary transmitter-receiver pair in the United States[3]. For this single-link example, we assume that the transmitter is operating at maximum height (30 meters) and maximum power (4 Watts EIRP[4]) on all available channels[5]. When not otherwise specified, these are the default secondary characteristics used throughout this paper.

We also assume that all TV transmitters are transmitting omnidirectionally at their maximum power and HAAT. Furthermore, we assume that our secondary receiver is able to attenuate noise from adjacent TV channels by 50 dB. Thermal noise (TNP) is present at all locations on all channels.

We will specifically look at two factors: interference from the primary transmitters (which we call *pollution*) and the need to halt transmissions at locations near primary receivers (which we call *exclusions*).

In Figure 2.3 there are four color-coded maps of the continental United States which indicate the data rate available to secondaries[6]. Each map represents a hypothetical world. In Figure 2.3a, the secondaries are operating in a clean channel (i.e. there are no other transmitters on that channel) and are allowed to operate on all channels at all locations. Adding noise (pollution) from the TV towers but still allowing them to transmit anywhere leads to Figure 2.3b and similarly limiting their location-channel combination but pretending they can ignore TV signals leads to Figure 2.3c. Finally, combining these two effects (pollution and exclusions) gives us the true state in Figure 2.3d. It is clear from these figures that exclusions are more painful than pollution to secondaries.

In order to understand how this affects the people of the United States, we calculate a complementary cumulative distribution function (CCDF) sampled by population for each map[7]. These CCDFs, shown in Figure 2.3e, support our earlier conclusion that exclusions are more painful than pollution. Even if secondaries could ignore TV noise, they can do no better than the green line. However, removing exclusions (but keeping the noise) would get us to the red line, whereas a clean and completely available set of bands is represented by the cyan line.

Note that the the gap between the red and cyan lines (adding in noise with no exclusions) is greater than the gap between the blue and green (adding in noise with exclusions) lines since the former allow transmission in areas which by definition are close to TV towers and are therefore noisier, resulting in a lower average SNR for the secondary.

This section has concentrated on the exogenous effects on secondaries, showing that exclusions are a much more serious problem than noise from TV towers. In the following sections, we focus on the endogenous effects that secondaries face.

[3]Thus interference from other secondaries is not considered; it will be discussed further in Section 2.4.

[4][2] actually specifies a maximum transmit power of 1 Watts but allows for up to 6 dBi of directional antenna gain, thus resulting in a maximum of 4 Watts.

[5]The FCC specifically mention that the maximum total secondary power is fixed and is *not* a function of the number of channels used. For illustrative purposes, we ignore this.

[6]Details on map generation can be found in Appendix A.

[7]Details on CCDF calculations can be found in Appendix A.

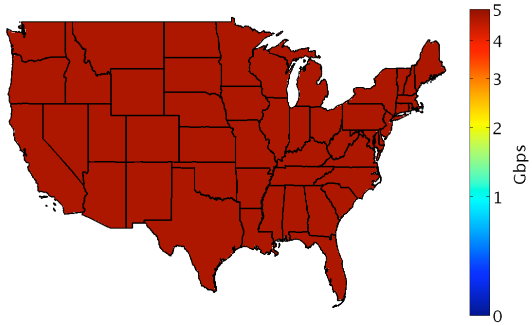

(a) No pollution, no exclusions

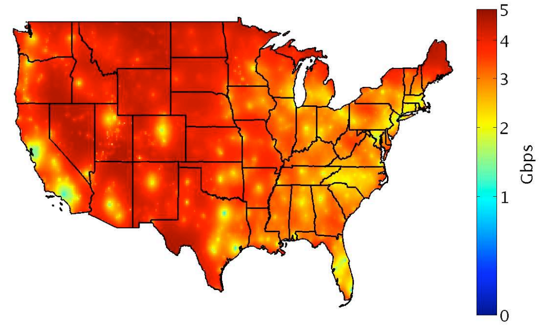

(b) Pollution, no exclusions

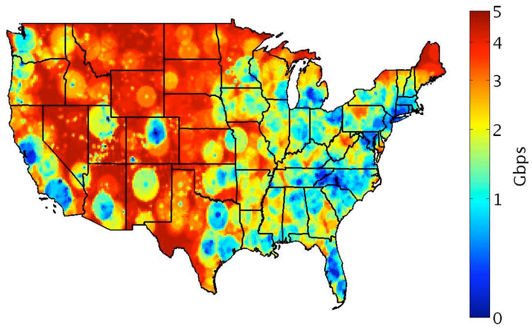

(c) No pollution, exclusions

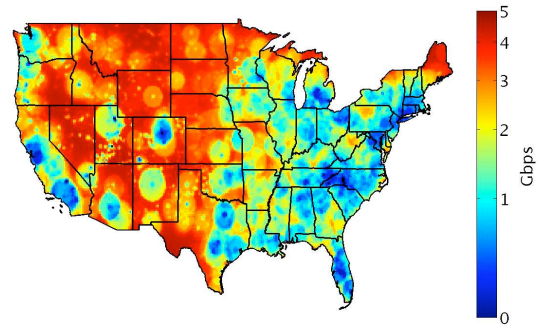

(d) Pollution, exclusions

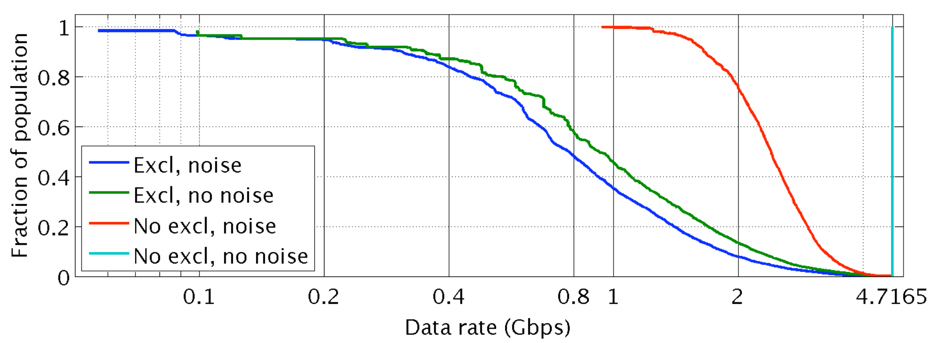

(e) CCDFs by population of maps (a)-(d)

Figure 2.3: Data rate available to a single secondary pair 1 km apart under various conditions.

2.3 Height vs. range: which matters more?

Height and transmission range are important parameters in any wireless system, including those operating in the TV whitespaces. In this section we will look at each factor separately and then show that range is dominant in most situations.

For reference, Figure 2.4 shows how a transmitted signal attenuates over distance in our propagation model for different transmitter heights[8]. Notice that height is inconsequential at some ranges.

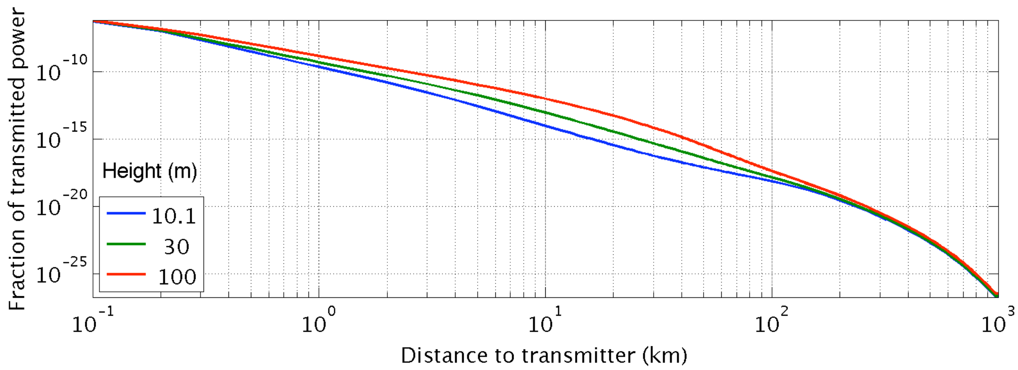

Figure 2.4: Signal strength as a function of distance from transmitter, height

Range

The transmission range of a communication link is the distance between the transmitter and the receiver. In Figure 2.5 we see just how dependent data rate is on the transmission range with the data rates collapsing in the 10-km case. This is due to the fact that a transmitted signal attenuates more with increasing distance, thus the SNR at the receiver is lower in the 10-km case.

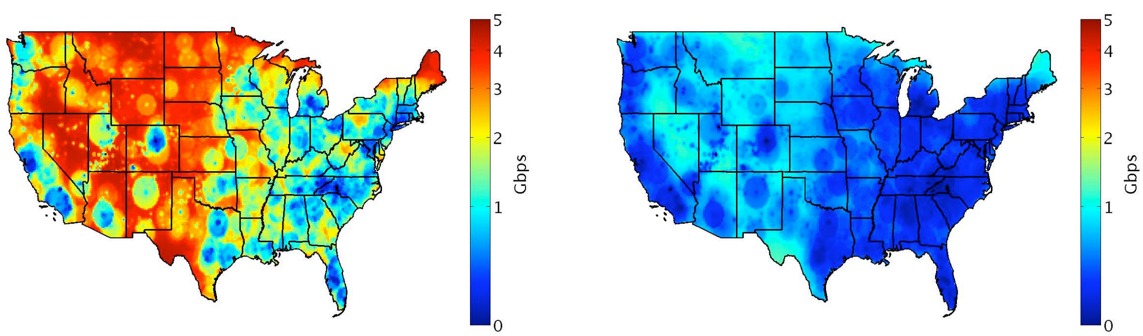

(a) 1km transmission distance

(b) 10km transmission distance

Figure 2.5: Data rate available to a single secondary pair with differing ranges.

Height

The height of a transmitter affects the propagation characteristics of its transmitted signal. Increasing the transmitter height increases the signal strength at any given location. However, we can see from Figures 2.4 and that this effect is minimal compared to that of range.

[8]All of our models use the ITU propagation model [10].

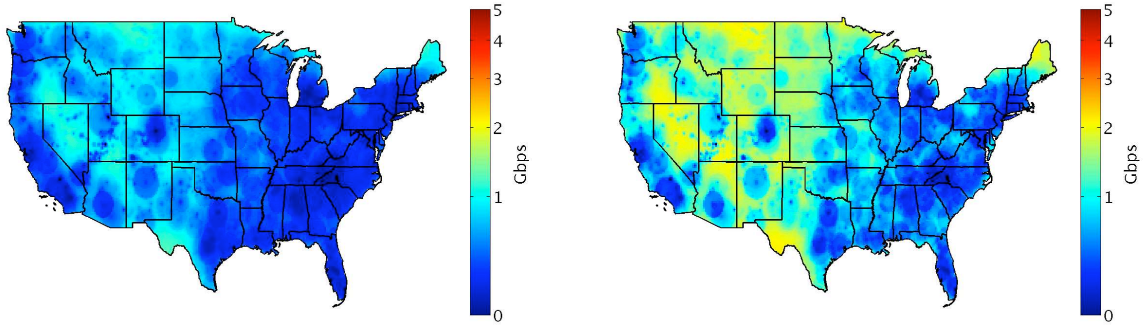

(a) 30m HAAT

(b) 100m HAAT

Figure 2.6: Data rate available to a single secondary pair with differing heights.

Comparison

To find the relative magnitude of these effects in the TV whitespaces, we created nationwide data for different (height, range) pairs[9]. For each pair, we found the rate of the average person and used that as a final data point, as shown in Figure 2.7. Again, the height boost only occurs at certain ranges since the signal strengths only differ at certain ranges. However, these ranges – approximately 5 km to 150 km – are precisely those that are likely to be used for long-range applications. In these cases, it is important to have and to exercise flexibility in transmitter height in order to maximize utility. Choosing the wrong height for a 50-km-range application may mean an order-of-magnitude decrease in utility.

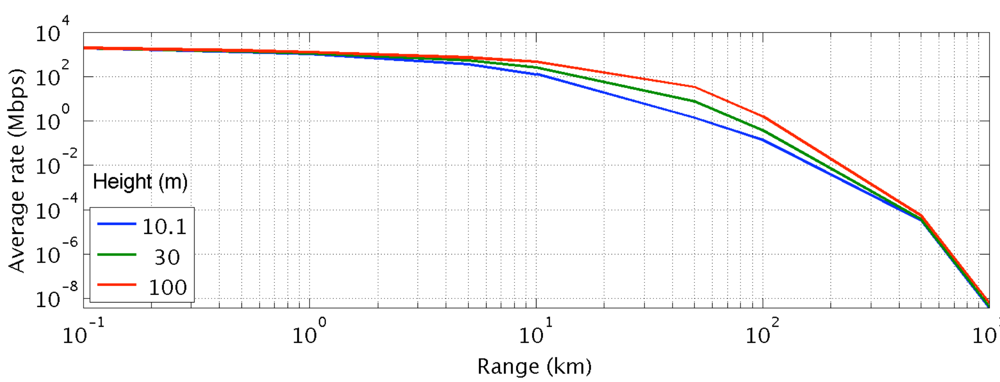

Figure 2.7: Nationwide average achievable secondary rate as a function of range and transmitter height

2.4 Self-interference: an unfortunate reality

The simple single-link model used in the previous sections promises incredibly high data rates because it ignores the effects of self-interference among secondary users. To a receiver, any unintended signal it receives

[9]This comparison isn't entirely fair since the FCC likely would not have chosen the same power limit if they allowed secondaries that were 100 meters tall.

is considered noise (also known as interference), even those from other secondary systems. This section implements a secondary sharing protocol to enhance the model and shows that the single-link rates are absurdly optimistic.

The Shannon capacity formula can be written more explicitly as

$$\text{rate} = \text{bandwidth} \cdot \log_2 \left(1 + \frac{\text{received signal power from transmitter of interest}}{\sum \text{received TV signal power} + \sum \text{received power from other secondaries}} \right)$$

Thus each secondary suffers a performance loss due to nearby secondaries that are concurrently transmitting. Some systems mitigate the effects of self-interference by using a medium access control (MAC) scheme in which neighbors take turns transmitting [59]. We will now consider such a system.

Specific model details are given in Appendix A.5. In essence, we ban all but one secondary transmission within a certain area called the *MAC footprint*. The *capacity per area* is then the transmitting link's rate divided by the MAC footprint. This is shown in Figure 2.8 where we again see a large difference in the achievable data rates based on the range.

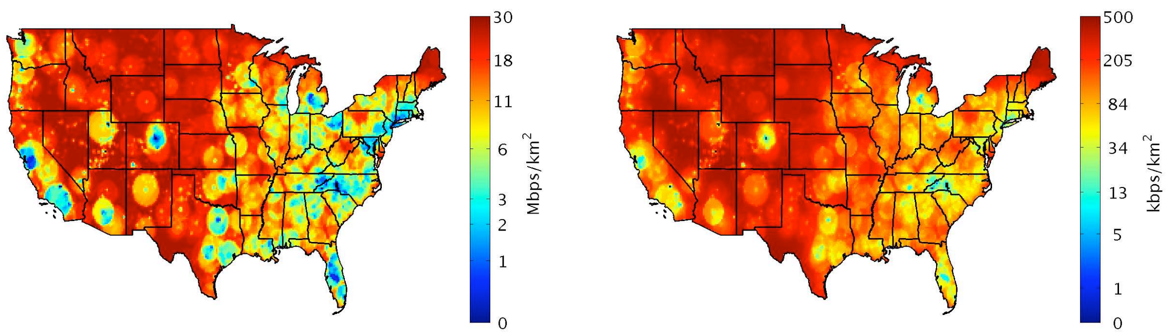

(a) 1km transmission distance

(b) 10km transmission distance

Figure 2.8: Capacity per area using MAC scheme

Further, we can consider the effective *capacity per person* by dividing the capacity per area by the local population density. The population density[10] is shown for reference in Figure 2.9 and the capacity per person is shown in Figure 2.10.

[10]Details on population data and calculations can be found in Appendix A.2.

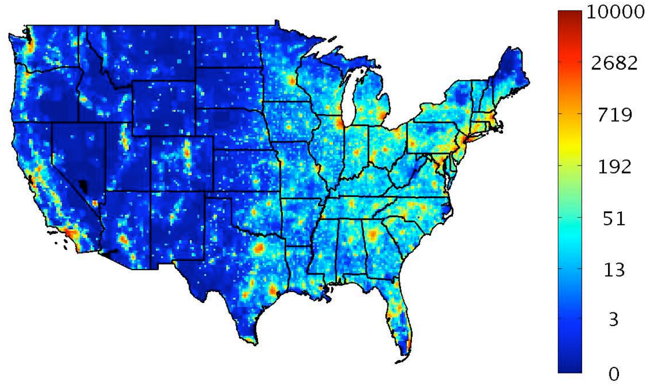

Figure 2.9: Population density (people per km^2) [4, 5]

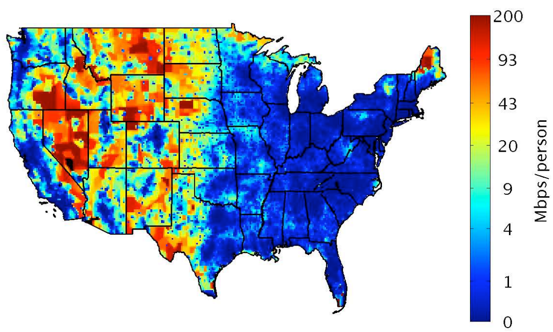

(a) 1km transmission distance

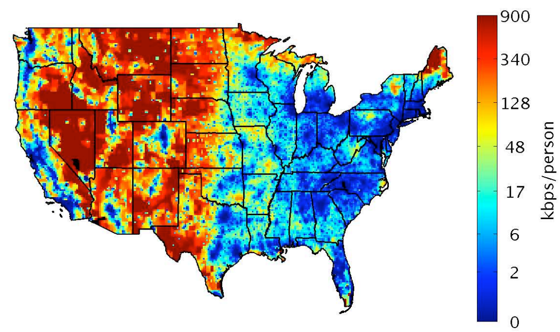

(b) 10km transmission distance

Figure 2.10: Capacity per person using MAC scheme

In Figure 2.11, we see the associated CCDFs for both the MAC model and the single-link model. The single-link capacity is clearly too optimistic as its rates greatly exceed those of the MAC model.

These maps and CCDFs also highlight the spread of achievable rates. In particular, we see extremely high rates in low-population areas and much lower rates in high-population areas. Recall from Section 2.1 that urban regions in general have fewer channels available for secondary use and the channels that are available are often noisier than in rural regions. Furthermore, the few resources available must now be split among many users whereas rural areas have many resources for few users. From these observations it seems that the TV whitespaces might be best suited for rural use. The following section addresses one of the potential flaws in this reasoning.

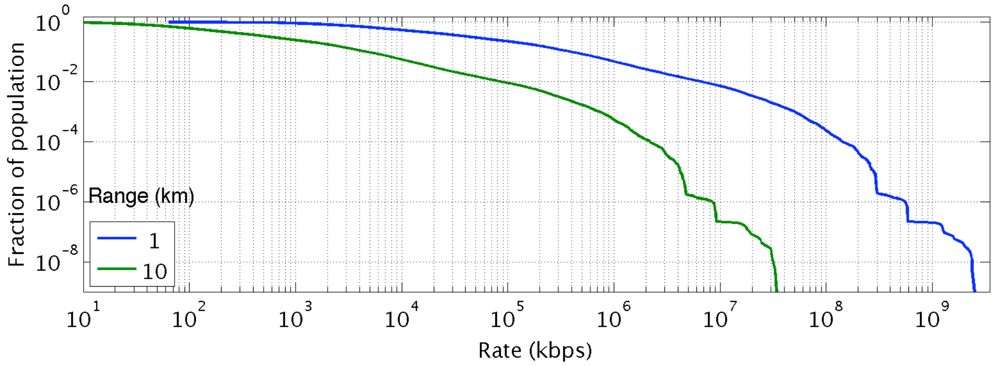

Figure 2.11: Distribution of rate per person in the MAC model with different transmitter ranges

2.5 Range tied to population: a qualitative shift

While some applications are immune to changes in population density, many are not. Consider a walkie-talkie example, an application which one may think of as fixed range. Suppose person A wishes to communicate with person B, where person B is chosen at random from the $p = 2000$ people nearest to person A (we tend to talk to our neighbors). In a city, person B is likely to live closer to person A than if they lived in the countryside and this is simply due to the relative population densities. Thus range is often tied to population density rather than being fixed at 1 km or 10 km. Specifically, we assume that one user in the pair is at the edge of the neighborhood containing p people and the other is at the center, as shown in Figure 2.12.

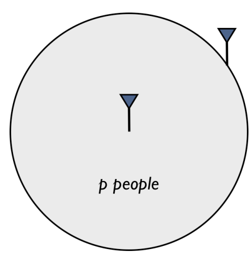

Figure 2.12: Single-link model where range is tied to population

Under this new model, low-population areas have much longer ranges than high-population areas. Figure 2.13 shows a qualitative shift: although rural areas generally have more available spectrum, their rates are typically lower than those in urban areas because of the longer ranges. In previous models where the range was constant, the situation was reversed: rural regions performed better than urban regions due to their extra spectrum.

Figure 2.14 shows how the size of the neighborhood, p people, changes the mean and median data rate in the single-link and MAC models. Intuitively the results make sense: the closer your communication partner, the higher your rate will be.

We discuss another important wireless model, the cellular model, in the following section. In this model, a service provider builds towers which will service non-overlapping regions. Each is sized to contain p people under the assumption that p customers are required to recoup the cost of providing the service.

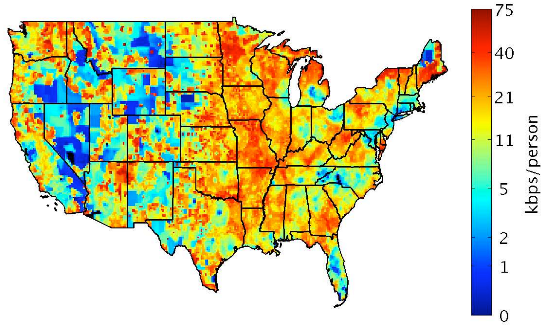

Figure 2.13: Capacity per person using MAC scheme with range tied to population density ($p = 2000$)

2.6 Cellular models

In addition to point-to-point links such as the walkie-talkie case, whitespace users may be interested in building a cellular architecture, for example as a wireless Internet service provider. We will show that the overall effect on data rates is similar to that of the MAC model.

While the MAC model captures the variety of ranges seen throughout the nation, it misses the variations in range on the smaller scale. In particular, a tower will not be communicating at a fixed range to its users, rather it will have a cell full of users with which it needs to communicate. This is illustrated in Figure 2.15a where the tower is at the center of the cell and the users are uniformly scattered within the cell[11].

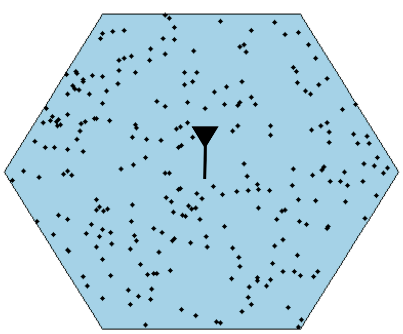

(a) User locations in one cell. The base station is located in the center of the cell.

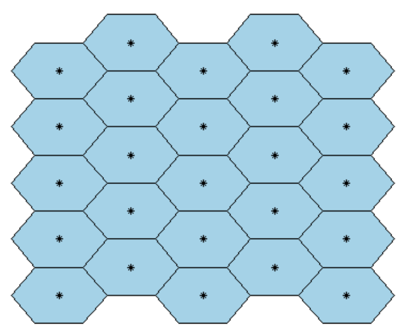

(b) Cell arrangement and tower locations.

Figure 2.15: Hexagon model illustrations

We consider a downlink-only model in which each tower transmits to each of the users in its cell in turn

[11]We cap the cell size to $\pi \cdot 100^2$ km^2.

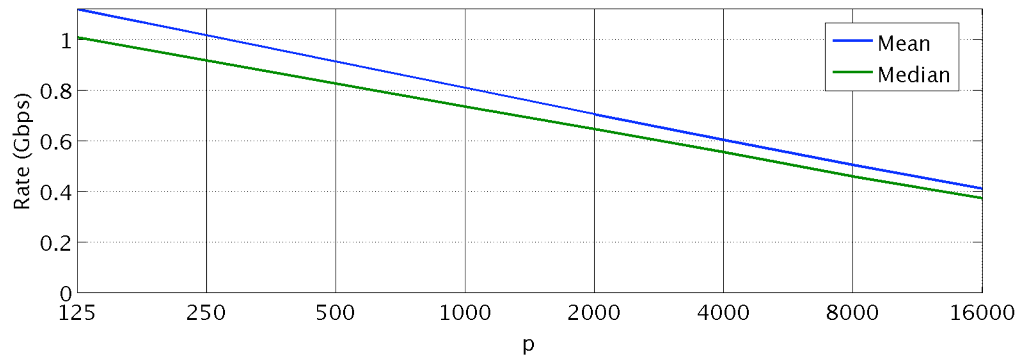

(a) Single-link model

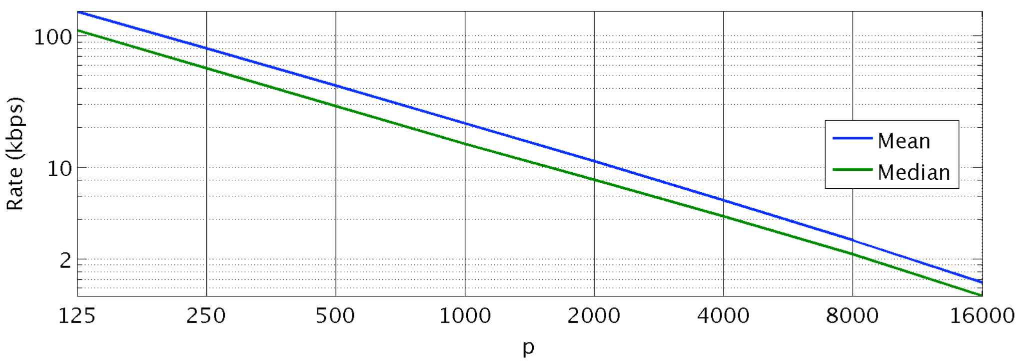

(b) MAC model

Figure 2.14: Effect of tying range to population in the single-link and MAC models.

and at a fixed power. Each receiver will experience interference from neighboring cell towers[12] as shown in Figure 2.15b.

The cellular model can be viewed as similar to the MAC model but with a fixed footprint. However, there is one important difference: secondaries in the MAC model can choose to accept any SNR[13], but the worst-case secondary at the edge of the cell almost always receives an SNR of approximately one-half regardless of the cell size and power used. This is because his predominant sources of noise are the two cell towers nearest him whose received signal strengths each match the strength of his desired signal. If P_r is the received power of each and N is the noise from the primary transmitters, we can write his SNR as

$$\mathrm{SNR} = \frac{P_r}{N + 2P_r} \approx \frac{1}{2} \quad \text{if } N \ll P_r$$

Following the spirit of the previous section, each cell is sized to fit p people[14], thus high-population areas

[12]All neighboring cells are assumed to be the same size which is based on the assumption that population is locally constant. As discussed in Appendix B, this is not correct in general.

[13]In cases where the pollution is too high, this may actually maximize the capacity per area.

[14]The idea behind this is that it takes p people to fund a tower's construction and maintenance. Back-of-the-envelope calculations suggest that $p = 2000$ is a reasonable number if we assume the following:

- Four people per household
- Households are willing to pay \$15/month for the service
- One-hundred percent market penetration
- \$50,000 per year is required to build and operate the tower

will have smaller cells which means both a shorter transmission distance but also greater self-interference. Conversely, low-population areas will have larger cells which means a longer transmission distance and less self-interference. We see this nationwide in Figure 2.16. As in the MAC model, we see rates in rural areas collapsing compared to rates in urban areas.

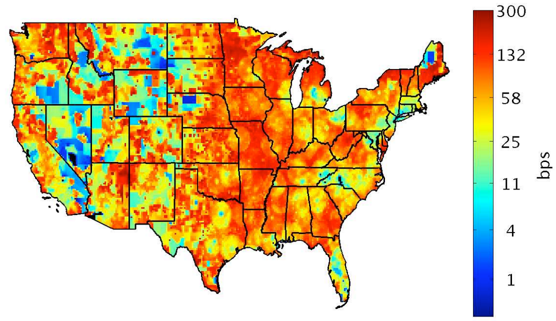

Figure 2.16: Hexagon model, $p = 2000$

The effect of changing the number of people per tower, as shown in Figure 2.17, is also similar to that in the MAC model.

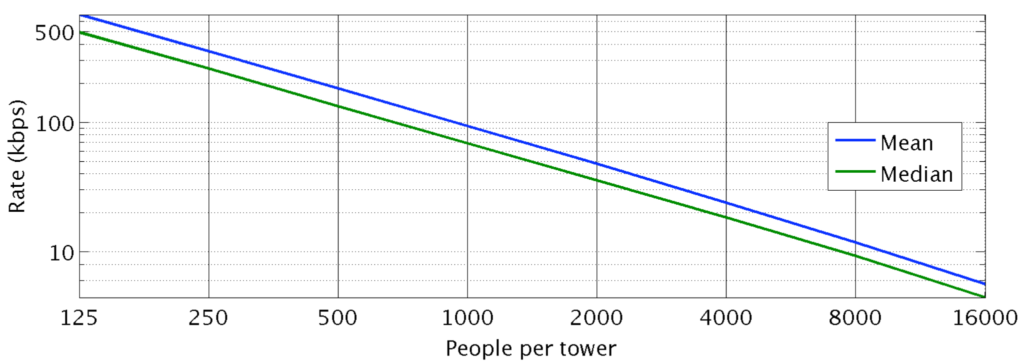

Figure 2.17: Effect of number of people per cell in the hexagonal-cell model

These figures give some insight into the secondary economy. In particular, they can be used to assess the viability of a particular business model and deployment strategy given the expected market penetration. One common option for reducing tower costs is to "piggy-back" on existing towers by paying their owner for the right to use their facilities rather than building a separate facility. This reduced cost would allow for smaller cells to be built, thus increasing the utility for the users of the system. In future work, we would like to assess the opportunity available in such an arrangement.

2.7 Conclusions

This section has taken a step toward quantifying the TV whitespaces opportunity and presenting some of its challenges. We began by showing that high-population areas have fewer channels than low-population areas which means that there will be less spectrum precisely where it is needed. Despite this effect, at least 87% of the US population has enough spectrum to operate a wireless router using the IEEE 802.11 standard which requires 22 MHz. However, secondary systems are subject to substantial noise from the TV towers ("pollution"), making this straightforward channel-count misleading. We suggest using achievable data rate rather than a spectrum count in order to more accurately quantify the whitespace opportunity. Using this metric and a simple single-link model, we then showed that although this pollution is troublesome, exclusions make the biggest difference in achievable secondary rates.

The secondaries' own system parameters also play a large part in determining their achievable rates. For example, we showed that an erroneous choice of tower height may result in a 50% decrease in rate on average. However, extending the transmission range can have a much greater effect, decreasing rates by up to two orders of magnitude with only a tenfold increase in range.

In addition to pollution, whitespace devices must also cope with interference from other secondaries. We first implemented a MAC sharing protocol which allowed secondaries to increase their individual rates via cooperation and time-sharing. The results from this model show that the single-link rates are overly optimistic.

The results from the MAC model also suggested that whitespace devices might be better suited to rural regions which have more spectrum and fewer people with whom to share it. However, we argued that, rather than being fixed, range is often tied to population density–even for a point-to-point link. Under this model, we saw a qualitative shift: urban areas now outperform rural areas despite their comparatively low amount of spectrum.

Finally, we developed another sharing method, the cellular model. Under this model, each secondary transmitter serves the nearest p users. Random user placement captures the variation in range on a smaller scale. Since cells are sized to hold p users, the variation in cell size captures the variation in range on a large scale.

Chapter 3

Changing the rules: is whitespace the right choice?

So far we have treated the FCC's regulations as if they were the only possibility rather than one choice taken from a set of potential rules. In this section we discuss two different families of rules, refarming TV channels and reducing protections for TV towers, and compare them to the FCC's choice.

3.1 Repurposing the channels: the traditional approach

Here, we consider the fundamental decision to even *have* TV whitespaces. Rather than create the TV whitespaces, the FCC could have simply reassigned TV stations to a smaller set of channels and opened up the remaining bands to unlicensed use. Alternatively, they could have refarmed some of the channels in this manner while opening up the remainder as TV whitespaces. The tradeoff for these two scenarios is shown in Figure 3.1 as the "no sharing" and "sharing" curves, respectively. The black crosshair marks the tradeoff point that the FCC chose.

In these calculations, we progressively remove channels[1] from TV use and give them over to exclusive "secondary" use. Each dot marks the "removal" of one channel, so the second dot indicates that one channel was refarmed, the third dot indicates two refarmed channels, etc. In the no-sharing model, channels still used for television are completely off-limits for secondaries. In the sharing model, secondaries may additionally transmit TV-inhabited channels using the current FCC rules for TV whitespaces[2]. Notice that the sharing model will do strictly better than the no-sharing model since the set of allowed (location, frequency) pairs under the sharing model is a strict superset of those in the no-sharing model.

There are two interesting features in this graph. First, note that the FCC is accepting a potential loss of over two channels per person on average just by allowing the existence of secondaries. This difference arises from the fact that a receiver near the limit of reception faces more noise with secondaries than without secondaries, thus fewer locations can receive TV. This issue will be discussed further in the following chapter.

Second, to achieve the level of service that the FCC has given to secondaries via refarming alone, they would have to relinquish an average of 3.5 channels per person in the $p = 2000$ cases, 9.5 channels in the

[1] Channels are heuristically ranked in order of greatest gain in a knapsack-like problem in which the value is the potential secondary-cum-primary rate and the weight is the number of TV viewers that will be lost.

[2] For simplicity we ignore adjacent-channel exclusions for the sharing model (the green line in Figure 3.1). Although this may seem a minor issue, we will see in the next section that adjacent-channel exclusions actually account for a large portion of lost secondary opportunity. We can see this already in the horizontal distance between the the FCC's tradeoff point (shown as a black crosshair) and the first green point. This green point represents the case where all channels are used as whitespace but with only cochannel exclusions; the FCC's tradeoff point is identical except that it also includes adjacent-channel exclusions.

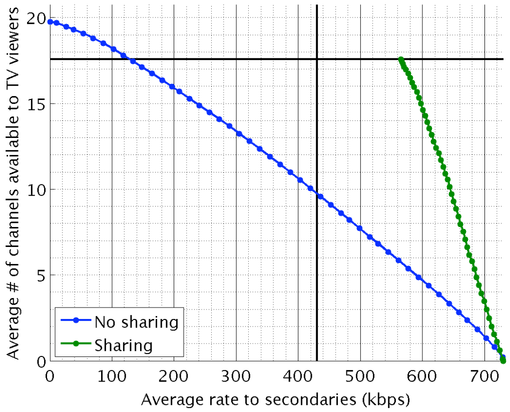

(a) 1km transmission distance, MAC model

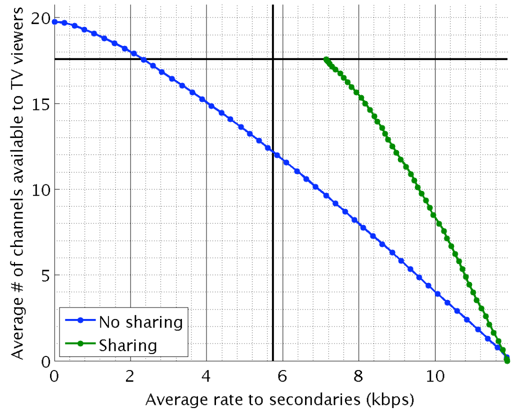

(b) 10km transmission distance, MAC model

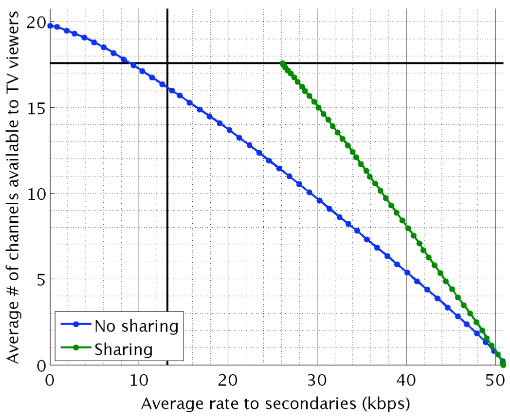

(c) $p = 2000$, MAC model

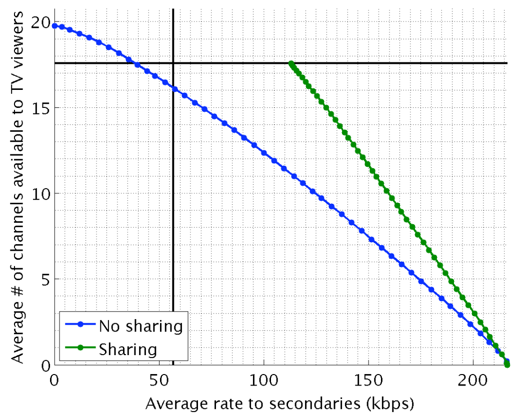

(d) $p = 2000$, hexagon model

Figure 3.1: Tradeoff between number of TV channels and secondary data rate

range = 1 km case, and 7.5 channels in the range = 10 km case. Thus whitespaces are more efficient than straightforward refarming for this level of secondary service.

Note that in these scenarios we have not accounted for where those TV stations will go should they be evicted from their current bands. It is not clear how this would affect these results. If the TV stations simply accept the loss or buy out another TV station, then nothing changes. However, if the TV stations choose to colonize new channels or locations, we would see a qualitative change: the curves of Figure 3.1 would likely be much steeper near the high-data-rate regions. In this scenario, the high population of TV towers in the remaining bands might cause so many exclusions that these bands would be all but useless to secondaries.

This section has shown that whitespaces appear to be better than typical channel reassignment given the FCC's desire to preserve most TV coverage areas. However the next section will show that whitespaces themselves can be designed in a variety of ways.

3.2 Making more whitespace: flexibility in the new approach

The selection of the protected radius (r_p) and the separation distance from the protected area ($r_n - r_p$) also represent a *choice* between secondaries and primaries. In this section we look at other options for these values and the tradeoff they present in terms of number of TV channels available and achievable data rate.

Here we consider a method of varying the value for r_p which is illustrated in Figure 3.2. This method was also developed in [40] but is recapitulated here for completeness.

Suppose that there exists only one TV tower, i.e. no other primaries nor secondaries. Due to signal attenuation and constant noise, the SNR for a receiver decreases as his distance to the TV tower increases. With digital television[3], the reception threshold is about 15 dB [60]; that is, when a receiver's SNR is below 15 dB, he can no longer watch TV. Thus there is already a naturally-defined maximum coverage area for each TV tower. We call this distance from the TV tower the 0 dB protected radius.

Clearly we cannot preserve this coverage area while allowing secondary operation. Any amount of additional noise, no matter how small, necessarily decreases the SNR and thus the coverage area. Therefore in order to add other transmitters on this channel it is necessary to sacrifice some of the coverage area, but how much?

We define an *eroded fade margin* to be the total amount, in decibels, by which the secondaries may decrease the primary receiver's SNR as compared to a clean channel SNR. For example, a margin of 3 dB allows secondaries to cause the same amount of noise as that found in a clean channel. At the same time, it decreases the coverage area to those areas which would have had an SNR of at least $15 + 3 = 18$ dB in a clean channel. Conceptually, increasing the fade margin decreases the protected area. We see this illustrated in Figure 3.2.

Note that the FCC's choice of protected radius does not in general directly correspond to any particular fade margin.

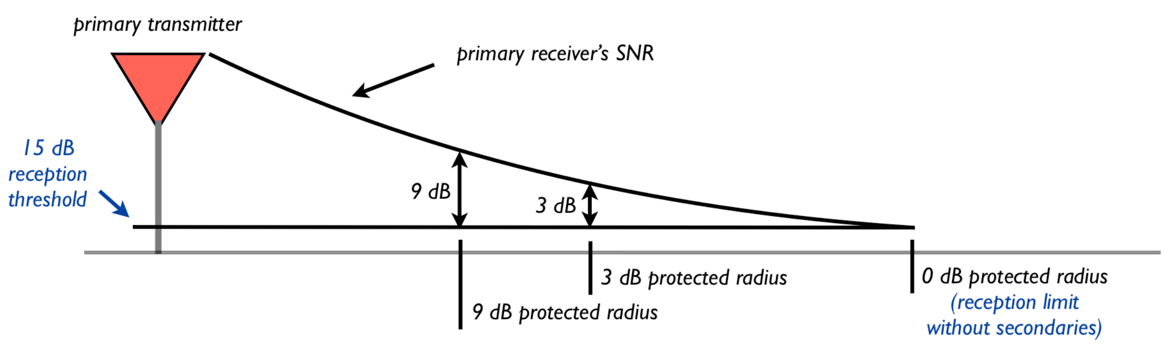

Figure 3.2: Fade margin illustration

In Figure 3.3 we see the effect on the secondaries of varying the fade margin. This helps us to visualize the impact of pollution and exclusions. Each point represents the median capacity (as sampled by population) of a single-user system with the corresponding exclusions.

Pollution alone appears to cause a great loss in capacity. However, recall that as in Figure 2.3e this is overstated since a complete lack of exclusions allows secondaries to use the extremely polluted areas they would otherwise be barred from.

[3] For simplicity, we assume that all current television stations are digital. In reality, many stations are still analog.

Notice though that the noise has more of an effect on the 10-km range cases; this is due to the fact that the secondary's SNR is much lower due to signal attenuation over a larger distance. This also explains the relatively small effect that noise has in the $p = 2000$ case: polluted areas are typically near population centers and this high population causes the range to shrink, thus the signal is often greater when the pollution is higher.

Among the two types of exclusions, cochannel and adjacent-channel, we see that the latter often has greater impact. Intuitively this makes sense because each tower has (typically) two adjacent channels, thus there are effectively twice as many towers to exclude on adjacent channels as on the same channel. Naturally as the fade margin increases, the overall effect of exclusions diminishes[4].

3.3 Conclusion: whitespaces best for minimal primary impact

We have examined two potential alternatives to the FCC's chosen regulations. The first alternative follows the traditional spectrum-reallocation method by refarming TV channels for unlicensed use, shown as the blue curves in Figure 3.4. The second, shown in red, creates whitespaces but varies the size of the protected region[5].

We clearly see in all four secondary-deployment scenarios that whitespaces represent the best tradeoff if few TV channels can be sacrificed. Once we are willing to lose about three channels, simple refarming generally gives better results than whitespaces.

Another option is to use a combination of refarming and whitespaces, shown by the green line in Figure 3.4. Refarmed channels are unlicensed but maintain a power limit and the remaining TV channels are used as whitespaces under the current FCC regulations[6]. This method always outperforms the simple channel-refarming scenario but of course they converge to the same point when all of the channels are given over to unlicensed use.

It is important to notice that even with relatively small excluded regions (i.e. the ends of the red and teal curves), whitespaces cannot achieve the same data rates as refarming. This is due to the inherent pollution from TV towers that secondaries must face. If a higher rate is desired, refarming is the only option.

We also see that this same tradeoff is very nearly approximated by varying the size of the protected region while ignoring adjacent-channel exclusions (shown in teal). We are unable to provide an explanation for this phenomenon at this time.

It is important to point out that the FCC's tradeoff point (the black crosshair) does not lie directly on the red line which represents varying the eroded fade margin using both cochannel and adjacent-channel exclusions. This is likely due to the fact that the FCC assumes the use of directional antennas for TV receivers but we do not.

These conclusions are important to take under advisement when setting the rules: they tell us when the whitespaces are the right choice and when they are not. However, we have failed to address two points:

1. Is the FCC's tradeoff point achievable under the existing regulations?
2. Do these curves accurately represent the limits of whitespace utility?

[4]The graphs in Figure 3.3 are misleading: as discussed in Section 2.2, shrinking the protected area indeed increases secondary utility but not uniformly. Areas near the TV tower are less valuable than those far away due to pollution, thus increasing the eroded fade margin has diminishing returns.

[5]This is identical to the eroded-fade-margin (EFM) method of the previous section. For each margin value we have calculated the TV reception maps in order to find the average number of channels remaining for TV viewing. This allows us to plot the number of viewable channels against the secondary utility in the EFM scenario.

[6]For simplicity we do not consider adjacent-channel exclusions in this calculation. In the future we would like to create a more sophisticated model which incorporates these effects.

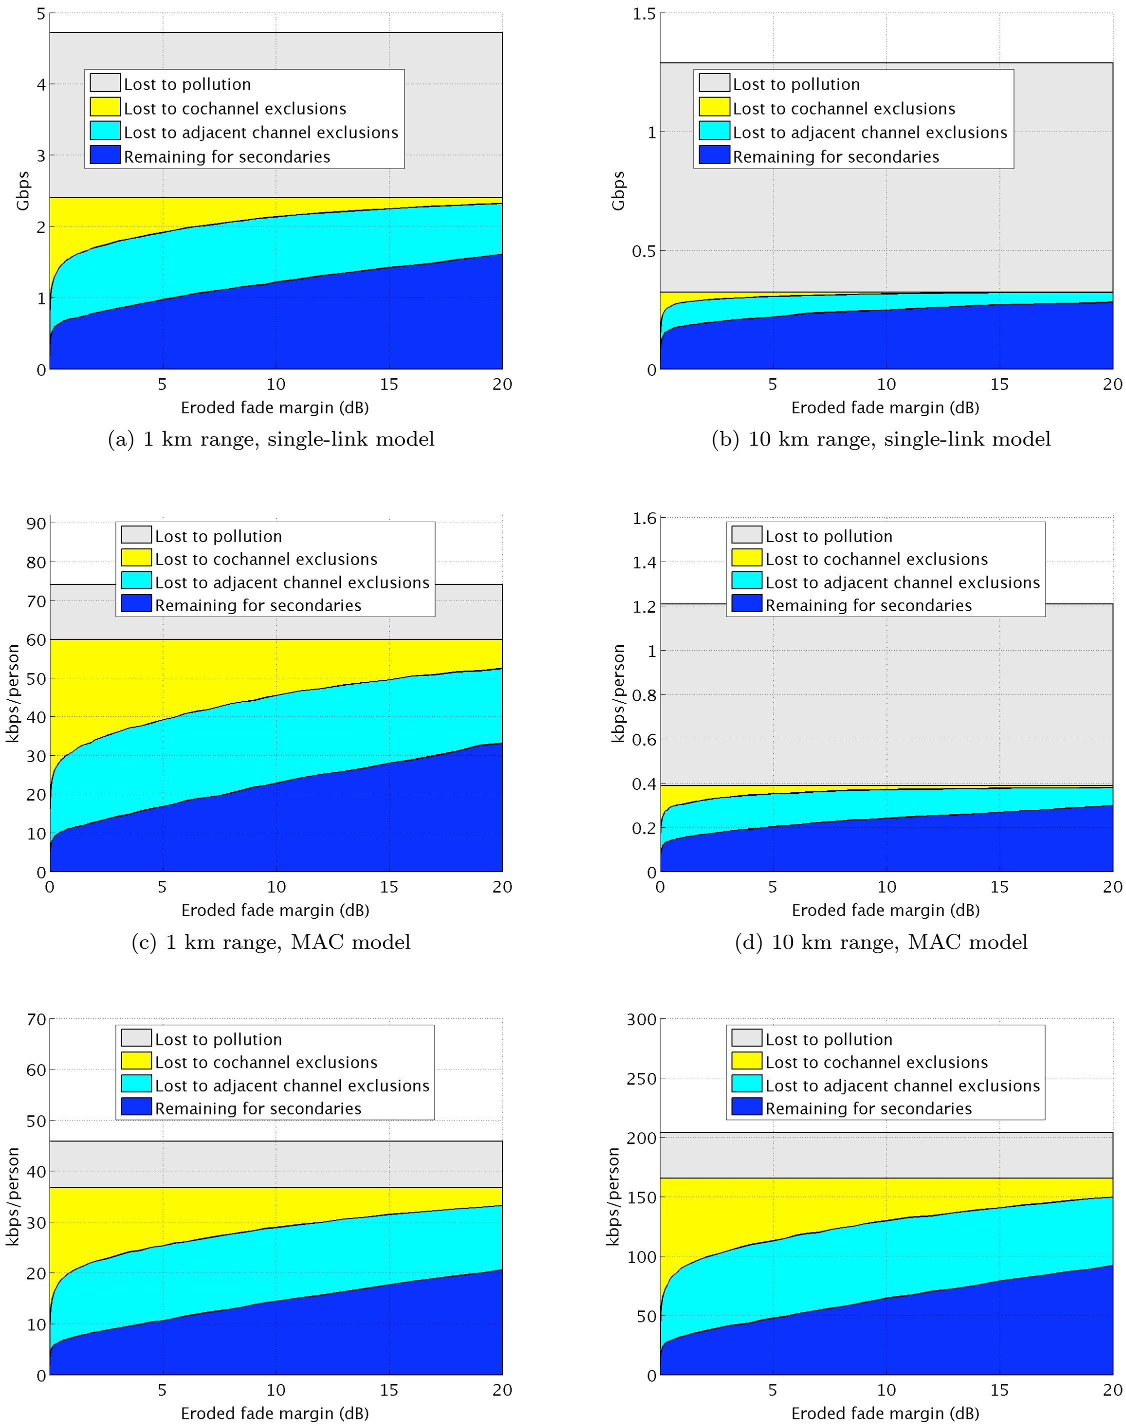

(a) 1 km range, single-link model

(b) 10 km range, single-link model

(c) 1 km range, MAC model

(d) 10 km range, MAC model

(e) $p = 2000$, MAC model

(f) $p = 2000$, hexagon model

Figure 3.3: Impact of TV noise and exclusions

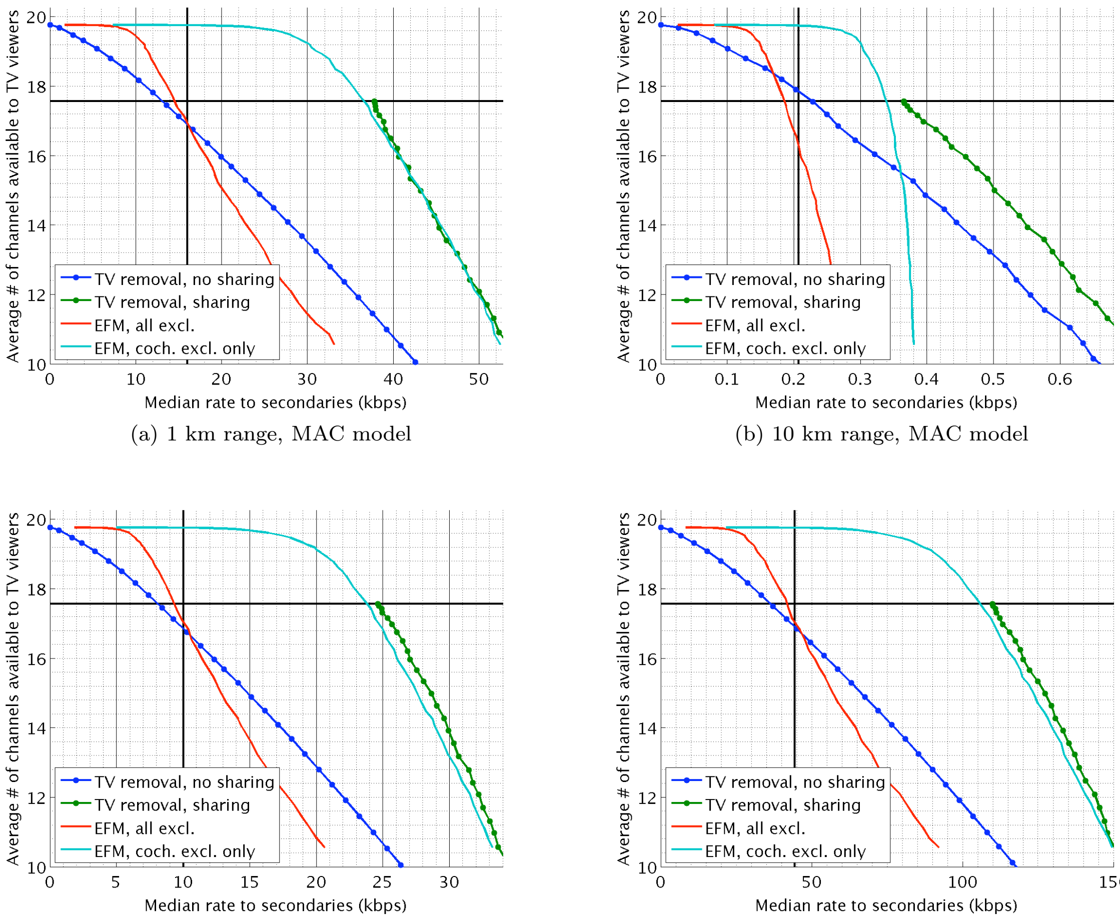

(a) 1 km range, MAC model

(b) 10 km range, MAC model

(c) $p = 2000$, MAC model

(d) $p = 2000$, hexagon model

Figure 3.4: Comparison of TV removal and eroded fade margin (EFM) schemes. The black crosshair represents the FCC's chosen tradeoff point.

We will show in the following section that the FCC's regulations fail to adequately protect primary receivers. We offer a remedy to this problem and then exploit this solution in order to improve whitespace utility.

Chapter 4

Strengthening the rules: ensuring primary protection while increasing secondary utility

In this section we show that due to interference aggregation the current FCC regulations may not be sufficient to provide the protection to primary receivers that they guarantee. Interference aggregation is the term which describes the cumulative effect of unwanted signals, in this case those from secondary transmitters, at the primary receiver. Since these regulations were designed to be safe in the presence of a single secondary transmitter[1], they fall short when multiple secondaries are operating. We will show the magnitude of this problem and suggest a remedy in the form of an enforced maximum power density, described in detail below.

We further show how this simple solution to the problem — made possible by the existence of databases — can actually be used to help solve another problem faced by the FCC: that of favoring one use case over another. We describe this tension below and offer a potential solution.

4.1 Good intentions and failures: interference aggregation

The regulations chosen by the FCC reflect a desire to use light-handed regulation for secondaries so as not to unnecessarily constrain their design. However, while the current rules do provide a remarkable amount of flexibility, it is also imperative that they maintain the quality of service that they guarantee to primaries.

[1] From the FCC's 2008 ruling in which they develop the separation distance requirements:

"In developing the table of separation distances, we believe it is desirable to minimize complexity for compliance. In this regard, we have balanced this goal of simplicity with the need to provide assurance that TV services will be adequately protected. Given that the power of fixed TVBDs will be limited to 4 watts EIRP, the most important variable in determining the separation distance between a particular TVBD and a TV station's protected contour is the height of the device's antenna above ground. For example, using the FCC curves in Section 73.699 of the rules and the D/U protection ratios specified above, we find that a transmit antenna at a height of 30 meters transmitting with 4 watts EIRP could cause co-channel interference to a TV receiver with an antenna 10 meters above ground at a distance of 14.4 kilometers and adjacent channel interference at 0.74 kilometers. For transmitting antennas at lesser heights, the FCC curves do not provide usable data, so the Okumura propagation model is applied. Using that same transmit antenna at less than 10 meters above ground, interference could be caused by a TVBD to a TV receiver at a distance of 8.0 kilometers to a co-channel TV station and 0.1 kilometers to an adjacent channel TV receiver. A similar calculation applied to a TVBD antenna at 3 meters above ground level calculates that interference can be avoided if separation distances of 6.0 kilometers and 0.1 kilometers are maintained for co and adjacent channel TV stations, respectively." [1, ¶181]

This is reaffirmed in the FCC's 2010 ruling:

"We affirm our decisions regarding the protection contours for TV stations. First, we decline to change the method that must be used to calculate TV station protected contours." [2, ¶21]

Due to the nature of the rules, the current limits do not provide the necessary level of protection. Interference from secondaries aggregates at primary receivers, potentially causing an outage. This is especially a concern in urban areas where the secondary device density is likely to be much higher.

To see the effect that interference aggregation has on a single TV tower's protected region, we create a toy world in which only one TV tower exists. This tower has a 500 meter height above average terrain (HAAT) and is transmitting on channel 21 with an ERP of 100 kW. We then place secondaries around the TV tower while respecting the FCC's rules. We assume that there is one secondary transmitter per 500 people and a population density of 379 people/km^2 (the median population density — by population — in the United States). Even though all secondary transmitters are obeying the FCC's rules, we see in Figure 4.1 that this actually shrinks the protected region's radius by over 10 km from 108.5 km to 97.4 km. This means that even those TV receivers which are well within the FCC's *protected region* may face too much interference to be able to receive TV.

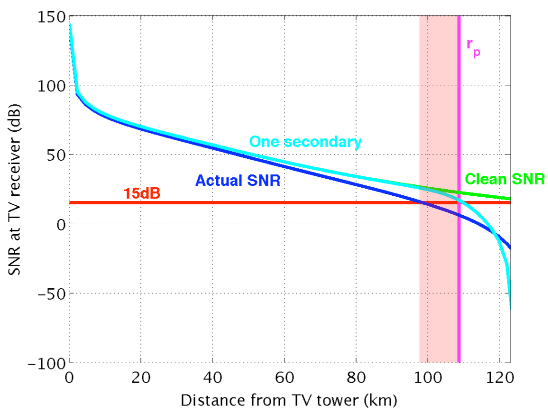

(a)

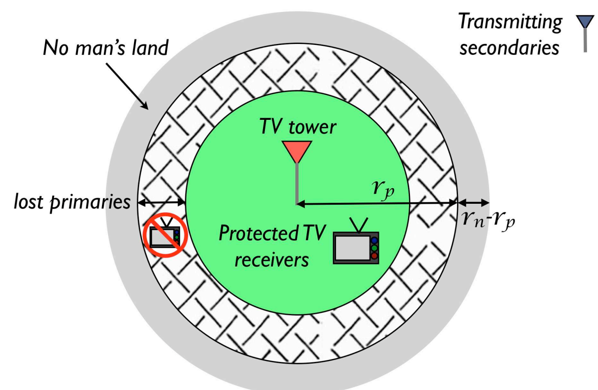

(b)

Figure 4.1: Lost area for one tower

To verify that this toy is not simply a pathological example, we perform a similar calculation for actual TV towers [3] using real population information [4, 5]. First, we find the magnitude of the maximum-strength TV signal for each pixel[2]. We then find the secondary interference level at each pixel given that the secondaries are obeying the FCC's rules and that there is one 1-Watt transmitter for every 40 people (think of wireless routers)[3]. These two pieces of information combine to form the signal-to-noise ratio (SNR) for TV receivers across the nation. We can then compare the areas with an SNR greater than 15 dB[4] to those areas inside of r_p. The difference is shown in Figure 4.2.

Note that we have shown only the locations that were guaranteed reception under the FCC's rules but have potentially lost reception; that is, we're omitting those locations that lost channels simply due to the FCC's ruling which were discussed in Section 3.1.

It is also important to understand that our nationwide model underestimates the amount of interference experienced at the TV receiver because of the scale: the large distance between pixels means that effect of local interference is severely understated. Since transmission powers attenuate very quickly with distance,

[2]Note that we ignore the possibility of both constructive and deconstructive interference from other TV signals. That is, we do not add TV signal strengths together, nor do we count the non-dominant TV signals as noise. The truth lies in one of these scenarios but the answer is different for each set of TV towers and we have not yet developed a way to determine which case we are in.

[3]We assume that all secondaries are transmitting all of the time. If you don't believe that that many devices would be active at the same time, consider the potential for hacking. Think of something like a wireless router for market penetration.

[4]We assume that 15 dB is the decodability threshold for TV signals [60]. Note that r_p is not explicitly designed by the FCC to have an SNR of 15 dB in the presence of secondaries.

local interference can actually account for a large portion of the total interference. Thus our estimates are much too conservative and need to be improved in future work.

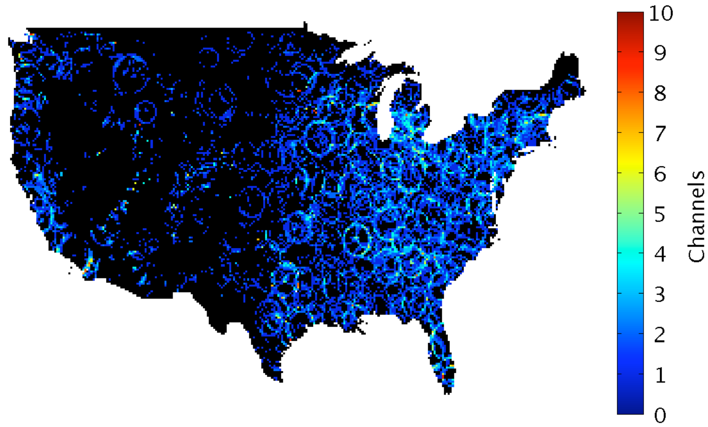

Figure 4.2: Number of channels at risk due to bad rules, $p = 40$

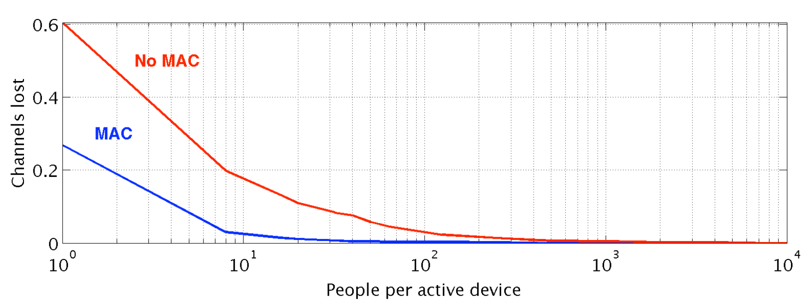

Figure 4.3: Conservative estimate of the average (by population) number of *protected* channels lost as we vary the popularity of secondary devices

We have discussed in previous sections that secondary devices might find it desirable to do some sort of self-regulated power density (e.g. using a MAC exclusion model or a minimum cell size). For example, two devices which are very near to one another would both prefer to take turns transmitting rather than have to deal with the other's noise. At some distance this is no longer true and the two devices will choose to transmit at the same time rather than take turns. This distance effectively sets the local power density in the MAC model.

However, this self-regulated power density is not sufficient to protect primaries, as shown in Figure 4.3 by the "MAC" line[5]. This stems from the fact that the self-regulated power density will not necessarily be the same as the "safe" power density.

The simplest solution[6] to this problem is to increase the separation distance $(r_n - r_p)$, but as Figure 4.4 shows, the population density in the United States varies too greatly for a one-size-fits-all rule to work well. In order to work, it would have to plan for the worst-case scenario. Since the difference between the 10th and 90th percentiles is two orders of magnitude, this would require serious compromise for much of the country.

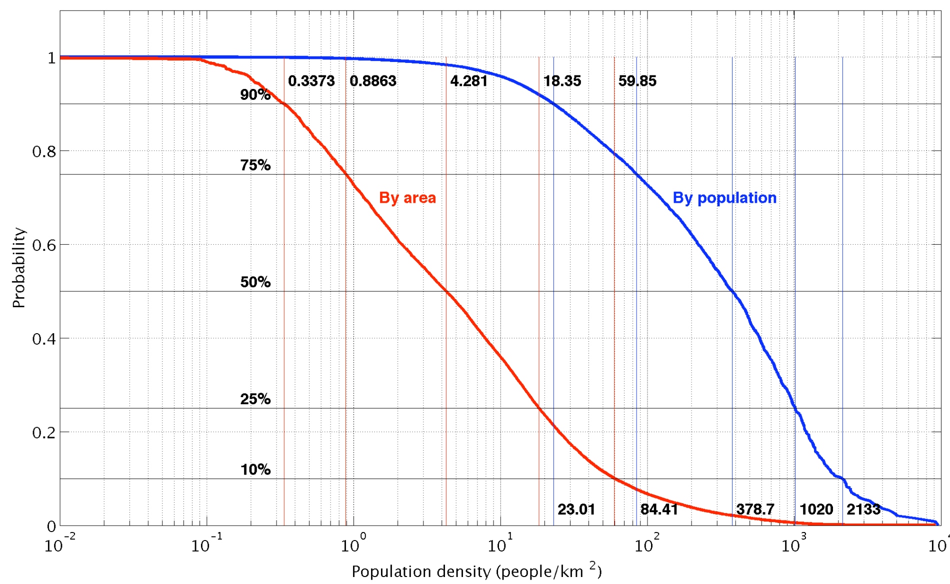

Figure 4.4: Population density of the United States, sampled by population and by area

It is impossible to anticipate the prevalence of whitespace devices, so we cannot design rules based on an expected device density. Instead, we must regulate a *power density*. We will illustrate this concept with a simple example. Consider two devices, A and B, which are co-located. Suppose device A is transmitting at the maximum safe power, P. When device B begins transmitting, their combined power can be no greater than P in order to maintain the power density and thus primary safety. As new secondaries turn on, nearby secondaries must adjust their power so that ultimately the rules are obeyed.

As discussed in Section 1.4.1.2, De Vany et al. [26] expressed very similar ideas with regards to how to control interference from secondaries. In their paper, they suggest that rights be defined in terms of an area in which the secondary is allowed to transmit. Furthermore, it shall be his responsibility to guarantee that his aggregate transmissions are below a certain threshold outside of his area. Most importantly, he suggested that anyone proposing to divide their transmission area (perhaps in order to sell part of it) would need to observe the original rules. That is, the worst-case scenario to neighboring areas remains the same regardless of whether or not the area is partitioned. In the extreme case where these property rights are partitioned into infinitesimally small pieces, we arrive at the notion of power density.

[5] In this case, we assume that devices must be separated by at least 200 meters in order to simultaneously transmit; that is, there is one active device per 0.0314 km^2. This is a reasonable distance for applications such as a short-range wireless service.

[6] It became apparent during the writing of this thesis that there is in fact another simple solution: choosing the separation distance $r_n - r_p$ for each tower individually (and perhaps dynamically) based on local characteristics rather than using the same separation distance for all towers. This would require minimal changes to the functionality of the databases and thus may be a more desirable solution to the problem of interference aggregation. The impact to secondaries of such a rule is a topic of future work.

There are two ways that the power density can be controlled:

1. A database gives a local power limit to a set of devices which control their powers collaboratively. Although harder to achieve, this decentralized solution will use spectrum more efficiently.
2. A database gives a power limit to individual devices. This has the advantage of minimizing device complexity and interoperability. However, it puts a large burden on the databases and relies heavily on propagation models [19].

Notice that both schemes require the various database providers to communicate with one another in near-real-time. However, this is not as imposing as it may seem: intercommunication between databases is already a requirement[7] due to the need to exchange time-varying registrations such as those for wireless microphones[8].

Both power-control methods are subject to the problem of an untruthful secondary which can either request more power than necessary or underreport its usage for malicious purposes. We do not consider these enforcement issues here. In our model, a region is given a power density and we assume that it perfectly obeys this mandate.

We took a rule of this nature and applied it to the United States in Figure 4.6a. We found the greatest safe uniform power density for each channel assuming that we are following the FCC's regulations regarding exclusions and that TV receivers at r_p can tolerate interference from secondaries that is the same as thermal noise (that is, we assume that a receiver at the FCC's r_p has an SNR of 18 dB in a clean channel and that 15 dB is required for TV reception [60]). In our model, each secondary is given power according to his footprint (in a traditional cellular system, this may be the coverage area of the cell). In particular,

$$\text{power} = \text{area of footprint} \times \text{power density}$$

First, we show in Figure 4.5 that using this power density is safe. Notice that we do lose a few channels when we compare the areas that receive TV to the FCC's protected regions. However, our models assumed that it was safe to erode 3dB despite the fact that the separation distance $(r_n - r_p)$ is not sufficient to guarantee this in some cases. Those stations for which this separation was inadequate will experience some outages in this model. Future work will incorporate this varying level of acceptable interference.

Figure 4.6b shows for comparison the rates for secondary devices under the FCC's current rules. While urban users receive roughly the same quality of service, rural users tend to do better because their larger cells allow them to use a power proportional to their area, thus helping to make up for their abysmally long ranges.

For a constant power density, increasing the area of a cell does not scale the allowed transmit power fast enough to offset the growing distance to the receiver. Consider an example of a cell which has radius r with a transmitter at the center communicating to a receiver at the edge of the cell, distance r away. Using a standard inverse-power-law pathloss model, the power necessary to maintain a constant rate to this receiver in a clean channel will increase with r^α where $\alpha > 2$. However, the transmit power is scaling only with respect to the cell area under our power density rule and thus is proportional to r^2. Therefore scaling the power based on area alone is not sufficient to maintain a set link quality in large cells. This suggests that we need to find another way of increasing the power available to rural users. We will address this problem in the next section by *safely* increasing the power density in rural areas.

[7] *"The Commission further required that, if multiple database administrators are authorized, the database adminstrators are to cooperate to develop a standardized process for sharing data on a daily basis, or more often, as appropriate, to ensure consistency in the records of protected facilities."* [2, ¶95]

[8] Events may register their venue as a temporary excluded region if they deem that the existing channels are not sufficient for their purposes.

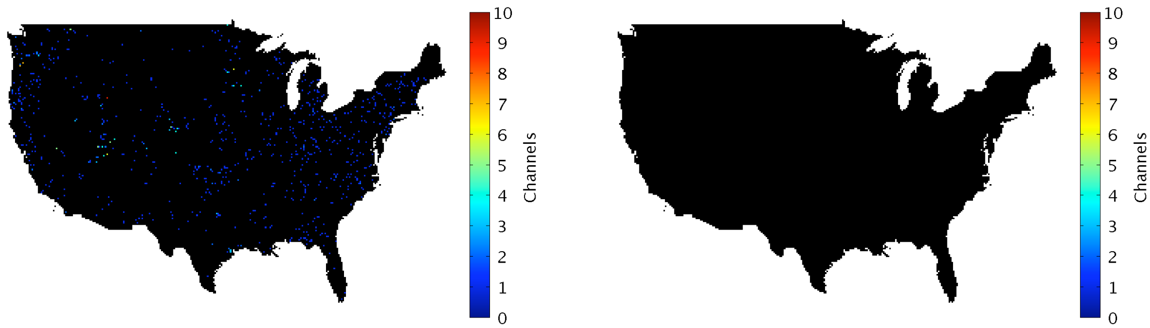

(a) FCC's protected region

(b) Regions which have an 18-dB SNR without the presence of secondaries

Figure 4.5: Number of TV channels potentially lost when using a uniform power density and two different protected regions.

4.2 Analysis: candidate rules

While it may seem straightforward to allow users far from TV towers to transmit with higher powers than those near TV towers, finding a *safe* and *fair* power scaling rule requires some thought. In this section we first show why the obvious solutions do not work and then offer a potential solution.

4.2.1 Why is power scaling even possible?

First, it is important to note that databases are the key piece of technology which enable us to realistically discuss power scaling schemes. Until now, unlicensed devices were certified for compliance with a uniform power limit–the only feasible option–and then deployed without any need for run-time adaptations. The ever-changing whitespaces (e.g. time-varying primary transmissions, event-specific wireless microphone exclusions) make such adaptability all but necessity.

Some devices will attain this adaptability via sensing, but we believe that many will make use of databases because they increase the *usable* whitespace opportunity [40, 41]. If we accept that these devices are already contacting the database daily[9] in order to receive a list of available channels, it is not outrageous to assume that they could also be assigned a varying maximum power limit.

In fact, Karimi of Ofcom has already suggested that this be part of the database capabilities [38]. Further, the Electronics Communications Commission of Europe has already suggested that rudimentary power scaling — but not power *density* scaling — may be advisable for maximum spectral efficiency and has provided a table of potential EIRP limits which vary as a function of the separation distance [19, §9.1].

4.2.2 Toy example: why doesn't straightforward power scaling work?

A naïve approach to power scaling would be to allow each secondary, regardless of its distance to r_p, to cause the same amount of noise at the protected TV receiver. We can quickly reject this scheme by realizing that

[9] *"The Commission required fixed and Mode II TV bands devices to re-check the database, at a minimum, on a daily basis to provide for timely protection of wireless microphones and other new or modified licensed facilities."* [2, ¶95]

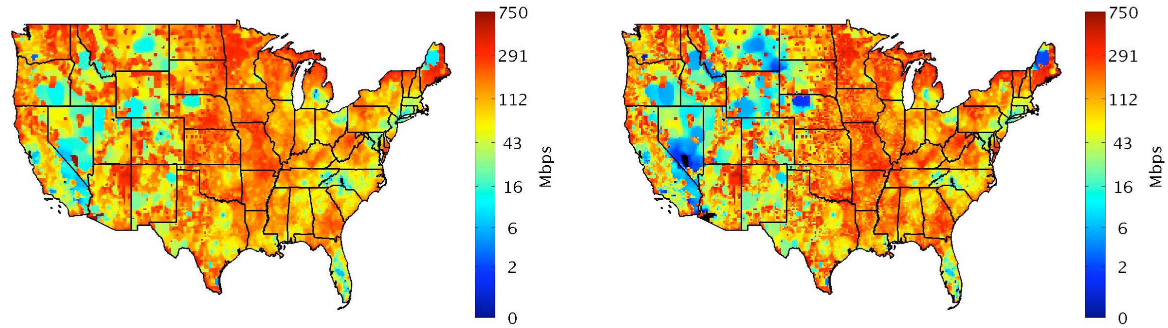

(a) Overly-cautious but simple power rule

(b) FCC's current rules

Figure 4.6: Rates available to secondaries under different power rules ($p = 2000$)

the eroded fade margin[10] would need to be split between a potentially very large number of secondaries, leaving each with only a small amount. Furthermore, this intuitively scales the power too quickly (it will scale as the inverse of the pathloss), thus giving regions far from the TV tower much more power then they need[11] while taking it away from nearby regions. Clearly we need a fairer way of dividing the margin. We will do this by looking at what each region "dreams" of having.

We refuse to allow secondaries to be wholly unrealistic in their "dreams" by pretending that they live in a world all by themselves. Instead, we take the viewpoint of a secondary who understands that r_n represents the boundary for secondary transmissions but that the value for r_n may be negotiated. Indeed, each person's dream is that they are living at the edge of the no-man's-land (r_n). In that way, they need not sacrifice power in order to allow anyone "in front of" them and they are allowed to use the maximum safe power. However, we cannot fulfill this dream for every user at the same time. We use the following toy model to explain why this is true.

Consider the toy one-dimensional[12] world illustrated in Figure 4.7. There is a primary transmitter, shown in red, and a primary receiver (i.e. television set) shown in green. Since this receiver is at the edge of r_p, optimality of the rules implies that he should be on the verge of losing reception when the secondary system is fully loaded. For this example, we will assume that this means that the aggregate secondary interference cannot be more than T, the amount of noise in a clean channel[13]. Thus the secondaries may collectively cause 3 dB of interference.

We use a standard theoretical model for the attenuation of a transmitted signal. In particular, if the transmitted power is P, then the signal strength distance r away is $P \cdot r^{-\alpha}$ where $\alpha > 1$ (in a two-dimensional world we would require $\alpha > 2$).

Consider a secondary transmitter which is at distance x from r_p, shown in blue in Figure 4.7. His dream is that he is the closest secondary transmitter to r_p, thus the power density is maximized for him. We will call his dream power density $P_{\text{dream}}(x)$. The aggregate interference must be less than or equal to T, and thus

[10]The eroded fade margin is defined in Section 3.2. Briefly, it defines the total allowed noise that may be caused at the TV receiver from all secondaries.

[11]Due to the saturated nature of the log function which describes their data rate, increasing the power has diminishing returns. Furthermore, devices which are in the interference-limited regime find that increasing their power is only marginally helpful: increasing their power necessarily increases that of their neighbors which has the net effect of decreasing the relative noise power from the TV.

[12]The forthcoming results can be generalized to two dimensions but for clarity we have used one dimension.

[13]Even a clean channel has "noise," called thermal noise.

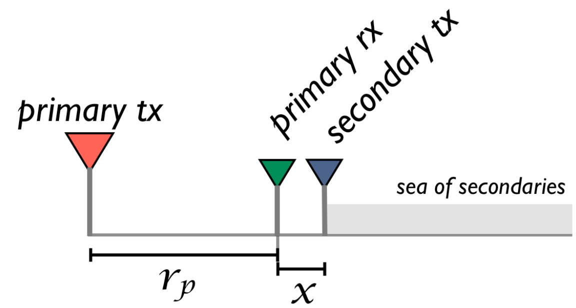

Figure 4.7: Toy one-dimensional world with one primary transmitter, one primary receiver at r_p, and a multitude of secondaries, the foremost of which is at $r_p + x$.

we have the following condition:

$$\begin{aligned} T &= P_{\text{dream}}(x) \int_x^\infty r^{-\alpha} dr \\ &= P_{\text{dream}}(x) \frac{x^{-\alpha+1}}{\alpha - 1} \\ \Longrightarrow \quad & P_{\text{dream}}(x) = T(\alpha - 1)x^{\alpha-1} \end{aligned}$$

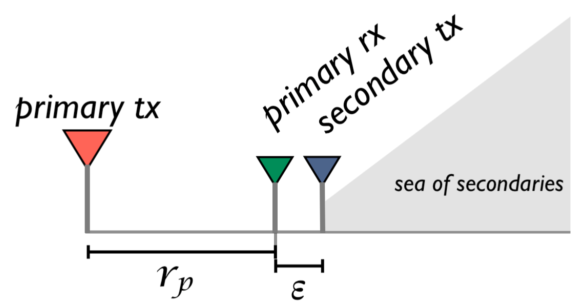

Figure 4.8: Dream power scaling for locations outside of $r_n = r_p + \epsilon$ (power density is shown on logarithmic scale).

While choosing this uniform power density will be safe under the secondary's assumption, it does not resolve the tension between the rural and urban users[14] since we must choose a particular x' to favor. To see why we can not give all their locations their dream power, suppose we chose $x' = \epsilon$ and let each secondary transmitter at location x_i use power density $P_{\text{dream}}(x_i)$ as shown in Figure 4.8. In this case, the interference at the primary receiver is unbounded:

$$\int_\epsilon^\infty P_{\text{dream}}(r) \cdot r^{-\alpha} \cdot dr \quad = \quad \int_\epsilon^\infty \left[T(\alpha - 1)r^{\alpha-1}\right] \cdot r^{-\alpha} \cdot dr$$

[14]For our purposes, we assume that TV towers are located at population centers. That is, urban users are those nearest to the TV tower while rural users are those furthest away.

$$
\begin{aligned}
&= T \cdot (\alpha - 1) \cdot \int_{\epsilon}^{\infty} \frac{1}{r} \cdot dr \qquad (4.1)\\
&= \infty
\end{aligned}
$$

The calculations above show that we cannot simultaneously give everyone everything that they want. Similar results were shown by Hoven [35] but he did not suggest a solution. We will now suggest a candidate rule to overcome this challenge in a way that is fair to urban and rural users alike.

4.2.3 Toy example: a candidate rule for fair power scaling

The previous section demonstrated why we cannot give all secondaries their "dream" power density simultaneously. We suggest the following fair compromise: if a secondary would have received rate R using his "dream" power density, we allow him to use the power necessary to achieve rate $\gamma \cdot R$ where $\gamma < 1$ is a tunable parameter which is used to ensure primary protection. Because of the logarithmic nature of the rate function, this results in a power density scaling which grows at a rate slower than $x^{\alpha-1}$, meaning that the integrand of Equation 4.1 is no longer $1/r$ and thus the interference at the primary receiver is bounded. This section will show these calculations in detail.

If a secondary would have received rate $R_{\text{dream}}(x)$ while using power $P_{\text{dream}}(x)$ in a clean channel, then under the compromise he will receive rate $R_{\text{new}}(x) = \gamma \cdot R_{\text{dream}}(x)$ by using power $P_{\text{new}}(x)$ where $0 \leq \gamma < 1$. The tunable parameter γ can be reduced until the SNR requirement for the primary receiver is met.

As we saw in previous sections, we can use the Shannon capacity formula to find the theoretical rate based on the bandwidth (we use a unit bandwidth in our calculations), the signal power, and the noise power [48]:

$$
\text{rate} = \text{bandwidth} \cdot \log_2 \left(1 + \frac{\text{signal power}}{\text{noise power}} \right)
$$

Our pathloss model dictates that signal power $= P \cdot d^{-\alpha}$ when P is the transmit power and d is the transmission distance. Notice that in order to calculate a rate, *we must assume a transmission distance, d.* We will come back to this problem later. Furthermore, in order to calculate the transmit power we need to multiply the power density by the footprint or "area" of the link, A. In our one-dimensional world we can think of area as length.

We assume incorrectly but for simplicity[15] that we are operating in a clean channel, thus noise power $= T$.

We will now show that adding the parameter γ does indeed cause the integral to converge. To begin, we calculate the old rate and corresponding new rate[16]:

$$
\begin{aligned}
R_{\text{dream}}(x, d, A) &= \log_2 \left(1 + \frac{A \cdot P_{\text{dream}}(x) d^{-\alpha}}{T} \right) \\
&\approx \log_2 \left(\frac{A \cdot P_{\text{dream}}(x) d^{-\alpha}}{T} \right)
\end{aligned}
$$

$$
\begin{aligned}
R_{\text{new}}(x, d, A, \gamma) &\approx \gamma \cdot \log_2 \left(\frac{A \cdot P_{\text{dream}}(x) d^{-\alpha}}{T} \right) \\
&= \log_2 \left(\frac{A^{\gamma} \cdot P_{\text{dream}}^{\gamma}(x) d^{-\alpha \cdot \gamma}}{T^{\gamma}} \right)
\end{aligned}
$$

[15] We would otherwise need to take into account not only the noise due to the primary but also noise from all other transmitting secondaries.

[16] We assume that we are not in the linear region of the $\log_2(\cdot)$ function.

$$= \log_2 \left(\frac{\left[A^\gamma \left(T^\gamma (\alpha-1)^\gamma x^{\gamma(\alpha-1)}\right) d^{\alpha(1-\gamma)}\right] d^{-\alpha}}{T^\gamma} \right)$$

$$\Longrightarrow \qquad P_{\text{new}}(x, d, \gamma) = A^\gamma \cdot \left(T(\alpha-1)^\gamma d^{\alpha(1-\gamma)}\right) \cdot x^{\gamma(1-\alpha)}$$

Notice that the power now scales with $x^{\gamma(1-\alpha)}$. This means that the integrand is now $r^{-(\gamma+(1-\gamma)\alpha)}$:

$$\int_\epsilon^\infty P_{\text{new}}(r, d, \gamma) r^{-\alpha} dr = A^\gamma \cdot T(\alpha-1)^\gamma d^{\alpha(1-\gamma)} \int_\epsilon^\infty r^{-(\gamma+(1-\gamma)\alpha)} dr$$

Thus we require $\gamma + (1-\gamma)\alpha > 1$ for convergence. We easily verify this by noting that $\alpha > 1$, $\gamma > 0$, and

$$\gamma + (1-\gamma)\alpha > 1 \Leftrightarrow (1-\gamma)\alpha > 1 - \gamma$$

Further, we can easily see that the aggregate interference condition is met if and only if

$$A^\gamma \cdot \frac{(\alpha-1)^{-(1-\gamma)}}{1-\gamma} \cdot d^{\alpha(1-\gamma)} \cdot \epsilon^{-(\alpha-1)(1-\gamma)} \leq 1 \tag{4.2}$$

Using this equation we can find the optimal γ if we can determine A, d and ϵ. However, this still assumes that there are no secondaries closer than ϵ to r_p, which is analogous to the FCC's $r_n - r_p$. We see how γ depends on ϵ in Figure 4.9. We can see already that the optimal choice of ϵ (a.k.a. $r_n - r_p$) is very dependent on the transmission distance, d.

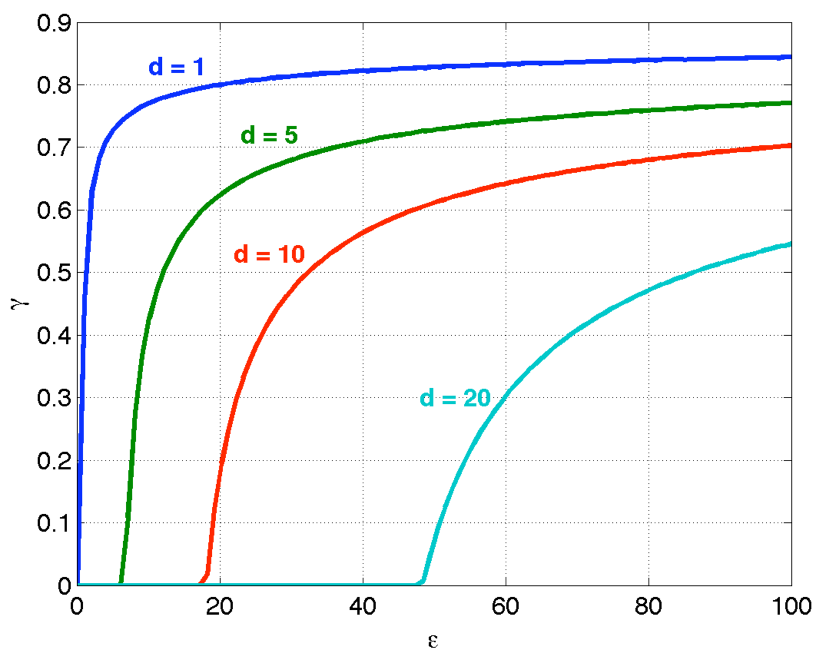

Figure 4.9: Optimal γ as a function of ϵ from Equation 4.2.

Of course, secondaries closer than ϵ to r_p will not be able to transmit at all on this channel and thus have zero *whitespace* rate. This may seem like a large sacrifice when viewed as a percentage loss, but on an absolute scale it is rather small if we choose ϵ carefully. These excluded secondaries would have a low rate anyway since γ would have had to be almost zero to include them and therefore we have not taken much away from them. Furthermore, devices always have the option of using the ISM bands instead of or in addition to the TV whitespaces. Finally, almost all people have access to at least two channels under the FCC's current regulations, implying they may be able to use another whitespace channel as well.

In our model, we choose ϵ dynamically (for each tower[17]) using a threshold β as described in Equation 4.3: secondaries whose dream rate does not meet this threshold are not allowed to transmit and secondaries who exceed the threshold are allowed to transmit.

$$\begin{aligned} \epsilon &= \arg\max_x x \\ &\quad \text{such that } P_{\text{dream}}(x) \leq \beta \end{aligned} \tag{4.3}$$

Thus we have obtained a (β, γ) approximately-optimal rule: the achieved rate for each secondary differs by no more than a factor of γ or the amount β from the dream rate. We now look at the effect of this rule on a nationwide scale.

4.3 Nationwide evaluation of candidate rule

In this section we will see the results of applying the candidate rule from the previous section to the United States. We will first describe the two models (cellular and hotspot) that we used to test our rule, then we will go on to contrast the two models in terms of dream powers and then show their realistic rates achieved. Finally, we'll show that TV reception is preserved under both models regardless of the device density.

The power rule developed in Section 4.2.3 is dependent on three variables: transmission range, the separation distance $\epsilon = r_n - r_p$ (set via parameter β as described in Section 4.2.3), and the rate-scaling factor γ. We choose the transmission range via the model as described in Section 4.3.1 below. We set the utility threshold $\beta = 0.5$ bits/second/Hertz since even this dreamed-of rate is not very useful. Finally, we maximize γ for each tower subject to the primary's interference constraint.

The interested reader is highly encouraged to review additional details regarding the generation of the following data which can be found in Appendix A.6 .

4.3.1 Two models: cellular and hotspot

Throughout this section we will consider two usage cases: the cellular case and the "hotspot" case. We define them here and henceforth refer to them by name only.

The cellular case has previously been described in Section 2.6. It essentially consists of a hexagonal cell with a transmitting (we consider downlink only) tower at the center and receivers (i.e. users) scattered uniformly throughout the cell. This is the typical cellular model. This is shown in Figure 4.10b where gray represents the footprint of the cell and green represents the support for the distribution of the users. For simplicity, we use a single user's range to calculate the rules but evaluate the rules based on all of the points in the cell. In particular, we consider a point approximately half-way out in the cell.

The hotspot model differs from the cellular model only in that users are clustered near the center of the cell, as shown in Figure 4.10b. Regardless of cell size, hotspot users are no further than 100 meters from their base station. One can envision a coffee shop scenario: there is one coffee shop on each block (hence the footprint size is one block) but each shop's wireless Internet signal does not extend beyond its walls. Again, we use a single user's range to calculate the rules but evaluate using all users. This range is 100 meters.

[17] This is done independently for each TV tower and the results are later synthesized into a nationwide map. Further details can be found in Appendix A.6.

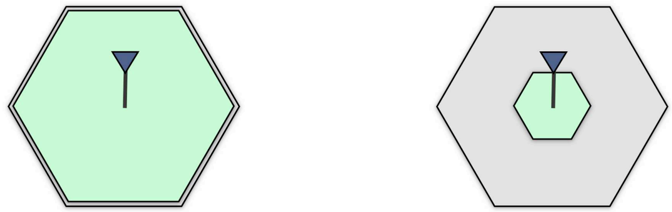

(a) Cellular model

(b) Hotspot model

Figure 4.10: Cellular vs. hotspot models: gray represents the footprint of the cell and green represents the area in which users may be. The transmitting secondary tower is located at the center of the cell. Cellular users are scattered throughout the entire cell and hotspot users are clustered near the center of the cell.

Under both models, each receiver will have a different potential rate which depends on his location in the cell, so in the spirit of fairness we assume that time is shared unequally among users so that each user gets the same overall rate regardless of his position in the cell. This is actually the harmonic mean of the potential rates. We call this the *fair rate* and it is our primary metric for evaluating cellular data rates.

As before, we set the cell footprint to be inversely proportional to the local population. We use the value $p = 2000$ throughout when not otherwise mentioned.

Again, many details have been omitted for the sake of brevity and readability. Please see Appendix A.6 for full details.

4.3.2 Powers and rates

In this section we will first show the "dream" powers and rates for each model (hotspot vs. cellular) and then we will show the results using the power scaling rule suggested in Section 4.2.3.

The average (across whitespace channels) dream power[18] is shown in Figure 4.11. Note that the dream power is the same for both the cellular and the hotspot models because it does not make any assumptions about the usage case. As expected, we see the dream power increasing as we go further and further from TV towers. Also, we see a high degree of variation in the dream powers. This further makes the case that a uniform power density would make a lot of places (and people) very unhappy.

We have applied a form of adjacent channel exclusions: we assume that adjacent-channel receivers can attenuate by 50dB[19] and then use the lowest of the cochannel and (potentially) two adjacent channel safe powers.

[18] Actually, we've taken the average not including zero-power channels.

[19] This was inspired difference in separation distances for cochannel versus adjacent-channel protections, presumably because TV receivers can reject adjacent-channel interference by approximately 50 dB.

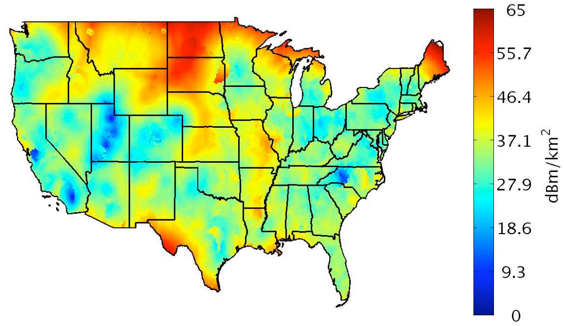

Figure 4.11: Dream power (same for cellular and hotspot models).

The rates resulting from using the dream power in each of the models are shown in 4.12. Our model incorporates both pollution and self-interference from nearby secondaries[20]. Note that this utilizes a new method for calculating secondary interference which is detailed in Appendix A.6.

As with the uniform power density, rural areas do quite well now because their large cells allow them to use a higher power which is almost enough to make up for their large transmission distances. The difference is that with the power scaling version of the power density rule, they now do even better than before because of the inverse relationship between population and whitespace channels seen in Section 2.1: fewer people means bigger cells but also fewer TV towers and hence a larger power.

We also notice that the hotspot model does extremely well. This is because the transmission range does not grow while the footprint does, thus it gets the best of both worlds.

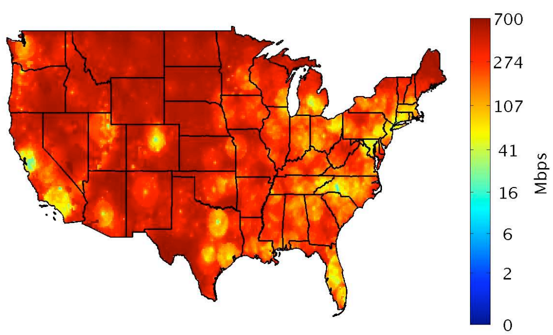

(a) Cellular model

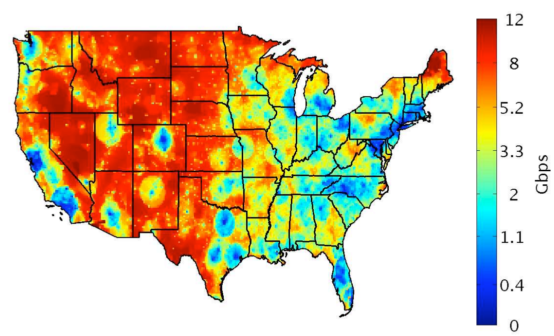

(b) Hotspot model

Figure 4.12: Dream rates

[20]Here, we no longer assume that nearby cells use the same power. Instead, we use the actual power densities.

These dream rates are unrealistic because the power density scales too quickly as we saw in Section 4.2.2. After we apply our candidate rule to the cellular model, we see that the resulting rates are about 20-30% lower than the dream rates, as shown in Figure 4.14. This means that most people are getting at least 70% of their dream rates. This is further corroborated by the CCDFs of Figure 4.16 where we see that at least 87% of people get at least 70% of their dream rate. The median person receives about 87% of his dream rate. And — most importantly — rural regions are doing about as well as urban regions.

The hotspot model naturally fares much better than the cellular model due to its constant short range. It loses only about 10% of its data rate from our candidate rule. Under this model, 96% of people get at least 70% of their dream rate. The median person receives about 92% of his dream rate. Rural regions do better than urban regions because in addition to generally having a greater power density their larger cell size allows them to use more total power.

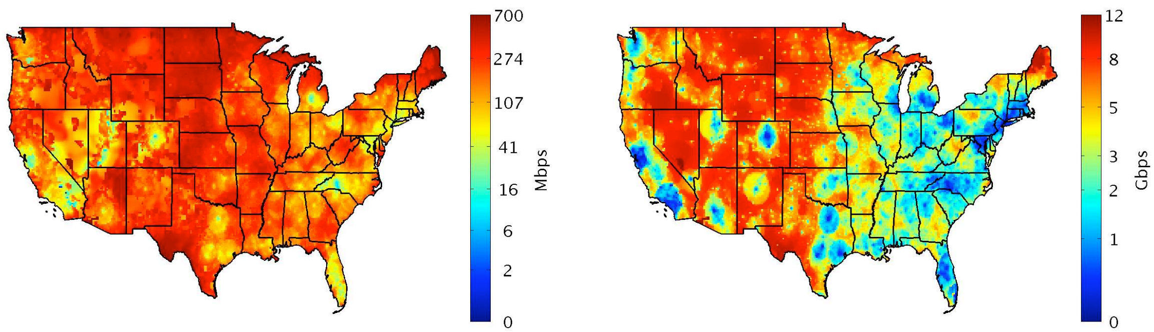

(a) Cellular model

(b) Hotspot model

Figure 4.13: Actual rates

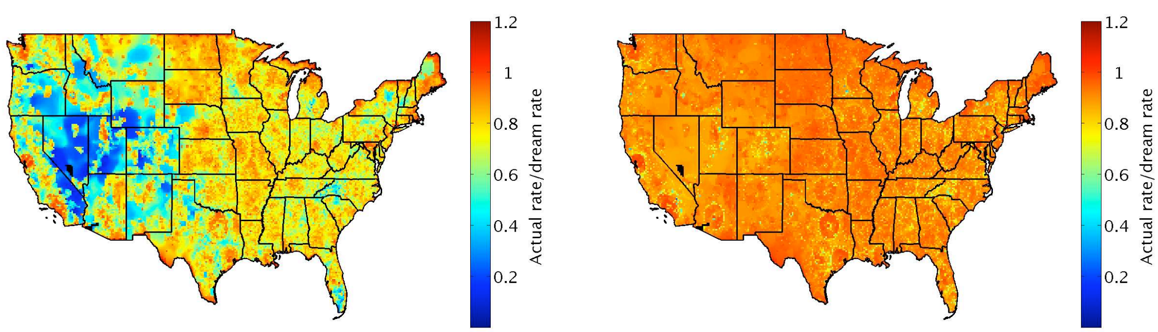

(a) Cellular model

(b) Hotspot model

Figure 4.14: Ratio of actual rate to dream rate

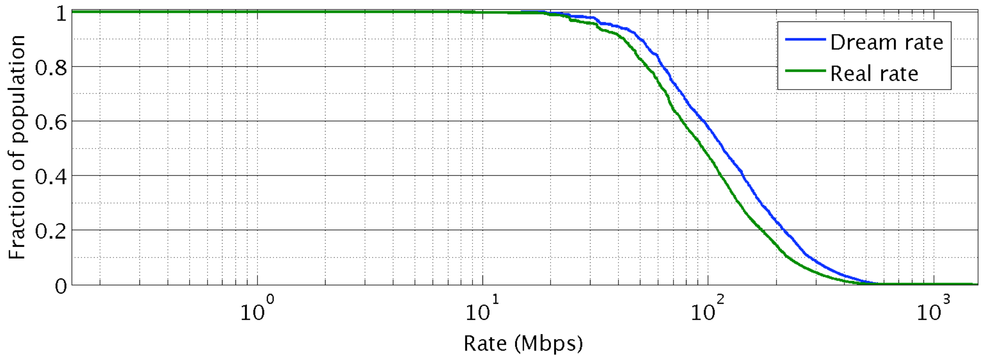

(a) Cellular model

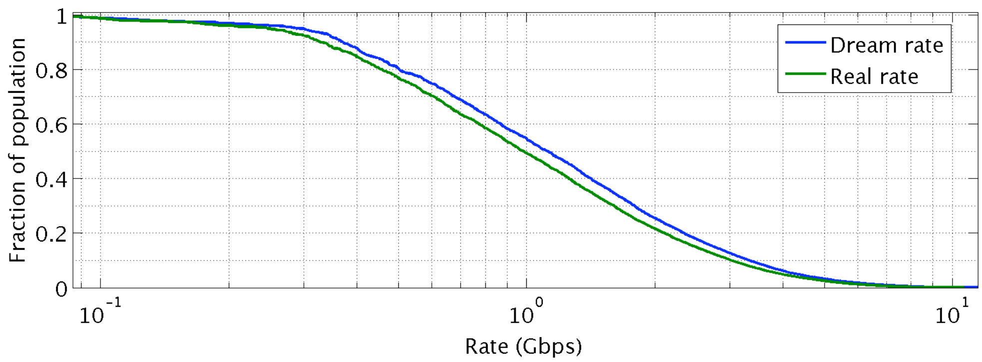

(b) Hotspot model

Figure 4.15: CCDFs by population of dream rates and actual rates

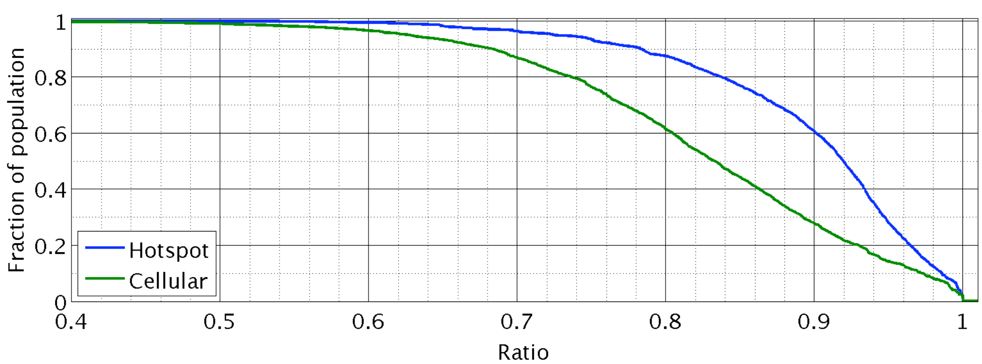

Figure 4.16: CCDF by population of ratio of actual rate to dream rate

As discussed in Section 2.17, we may have different "device popularity" levels, represented by a change in the value of p. We briefly show the effect of changing p on the average rate per person, shown in Figure 4.17. The effect seen here is two-fold: increasing the number of people in a cell not only increases its area (thus making a set rate harder to achieve) but also that same dismal rate is shared with even more people.

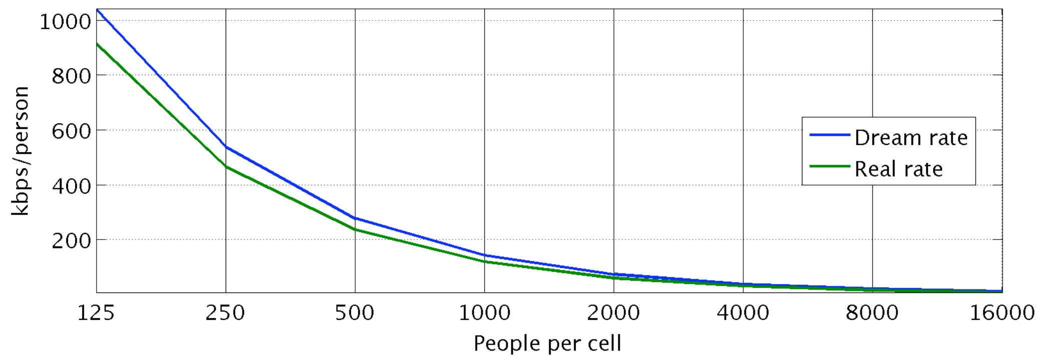

(a) Cellular model

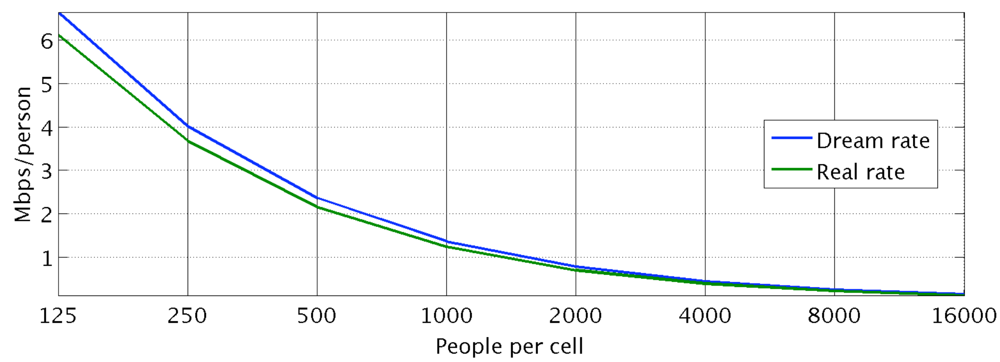

(b) Hotspot model

Figure 4.17: Average rate per *person* as a function of p, the number of people in the cell (the area of the cell is a function of p and the local population density)

A more interesting result shows up in Figure 4.18 where we plot the total rate per cell as a function of p. The hotspot model experiences an increase in utility because it has an increasing footprint and it does not have to cope with an increasing range. However, the cellular operator sees a downturn in his rates once his range grows too large. Notice however that his dream power is scaling fast enough for his dream rates to be generally increasing.

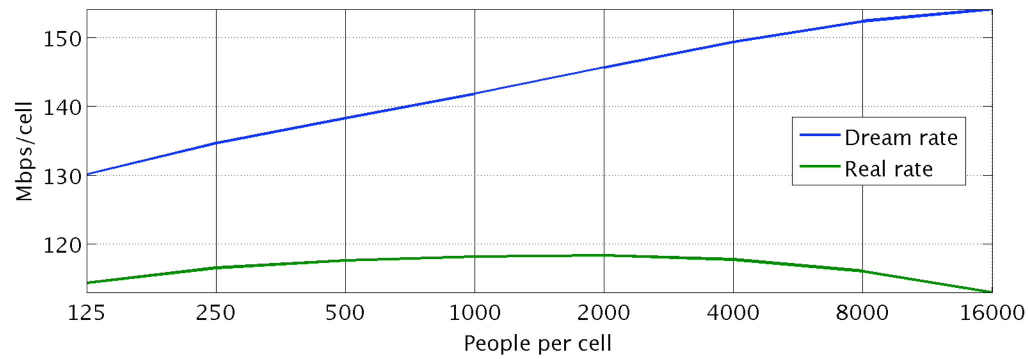

(a) Cellular model

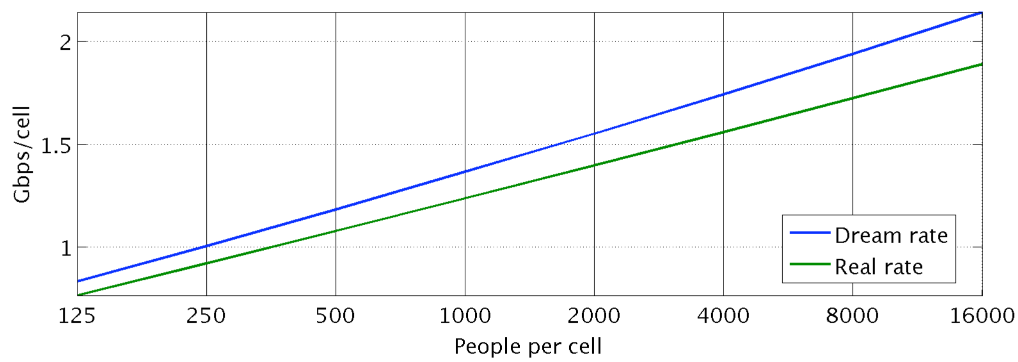

(b) Hotspot model

Figure 4.18: Average rate per *cell* as a function of p, the number of people in the cell (the area of the cell is a function of p and the local population density)

Finally, we do one last reality check to make sure that our power density rule is safe. This is done in the following section.

4.3.3 Reception for TVs preserved

In this section we simply provide proof that our scaling power density model continues to provide reception to TV receivers. We showed earlier that a uniform power density was safe and Figure 4.19 shows that indeed the scaled power density is also generally safe.

We note for future work the problem of investigating why the few at-risk regions exist.

4.4 Criticism of the candidate rule: tension between short-range and long-range applications

Although our candidate rule helps to mitigate the tension between the two usage cases (cellular and hotspot), this section explains why we have not completely overcome this problem. A complete solution is left for future work.

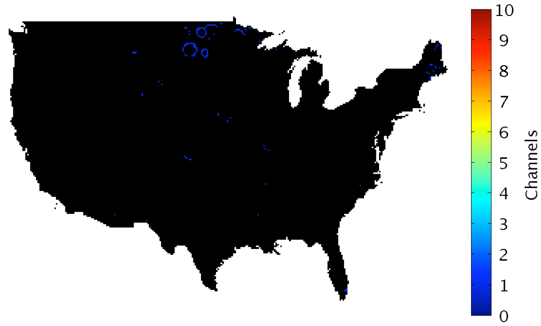

Figure 4.19: The FCC's protected regions are still protected under the power-density-scaling rule.

As we showed in Chapter 3, there are many ways in which to choose the protected region (r_p) and the size of the no-man's-land ($r_n - r_p$) for each TV tower. Each choice yields a power density limit, just as the FCC rules did. However, it is intuitively clear that users in urban areas want exactly the opposite of what rural users want. In particular, urban users who generally have a shorter transmission range and access to fewer channels than rural users (as shown in Section 2.1) will be willing to sacrifice some power in exchange for a greater coverage area. On the other hand, rural users are not nearly as hindered by exclusions and if acting purely for self-interest would lobby for a greater no-man's-land in exchange for greater transmission powers.

Unfortunately, our power-density-scaling rule requires us to assume a transmission range d and a footprint A when calculating the "dream" rate which in turn influenced the final power density level. When the mismatch between the expected deployed transmission ranges and footprints is high–in particular, high enough to cause a transition from the saturated region of the rate function (log) to the linear region or vice versa–the deployed devices will have a high amount of regret. Even having excess power is regrettable: receiving more power than necessary in one region necessarily means depriving another region.

It is not obvious how to resolve this problem in a fair way. We suspect that artificially introducing heterogeneity into the power density levels seen across channels at a given location may help: in this way, a device may be able to pick a channel whose power density has been optimized for his usage case.

4.5 Conclusions

At the beginning of this section we showed the potential consequences of the current FCC regulations. In particular, we showed that TV channels may be lost *in addition to those already lost by allowing the whitespaces* (see Chapter 3). This is an inevitable consequence of interference aggregation and it must be dealt with.

We propose a simple solution of using a power density in order to keep the aggregate interference at acceptable levels. We then show the rate available to secondaries under such a power density rule and compare it with the rate available under existing rules. Amazingly, this rule which is safer for primaries is also better for secondaries!

However, rural users are still suffering from ridiculously long ranges due to larger cells[21]. We seek to design a safe power scaling rule which mitigates this problem while maintaining quality of service for urban users

[21] This is a consequence of our "economic feasibility" model which was developed in Sections 2.5 and 2.6.

and settle on a candidate approximately-optimal power-density rule. Under this candidate rule we see that the achieved rates have roughly equalized across the country.

Finally, we acknowledged the shortcomings of our candidate rule suggested an as-yet untested solution. In particular, we were unable to find a rule which was agnostic of the secondary usage scenario. In future work we would like to find a rule which fairly navigates the tradeoff between the many different deployment scenarios.

Chapter 5

Conclusions and future work

This thesis first quantifies the opportunity to secondaries in terms of available spectrum and achievable data rates via simulations involving real-world TV assignment [3] and population data [4, 5]. Prior studies for the United States [40, 41] and the United Kingdom [43] have bounded the amount of available spectrum but fail to account for significant effects such as an increased noise level and self-interference among secondaries. These effects mean that the traditional bandwidth metric is inadequate since spectrum users are not interested purely in spectrum but rather in what they can do with it, which can be characterized by the achievable data rate. Our work incorporates this metric in order to better quantify the impact of the TV towers and associated regulatory decisions on the utility of whitespaces to secondaries. Through this metric, we verify that indeed exclusions are more painful to secondaries than noise.

We also develop models for both point-to-point and cellular systems in order to quantify the impact of self-interference among secondaries in the TV whitespaces. These models showed that the single-link model is inadequate to characterize the whitespace opportunity. In order to accurately assess the whitespaces it is important to understand the role of self-interference: while secondaries which are limited by self-interference are minimally affected by pollution, secondaries that are noise-limited find pollution the more costly aspect of operation.

Second, we improved on the work done by Mishra, et al., in [40] by comparing simple channel reassignment to the use of whitespaces. We conclude that whitespaces are indeed the best choice if only a few TV channels can be sacrificed. However, the pollution from primaries mars the whitespaces enough that additional channel reallocation may necessary if high secondary utility is needed.

Third, we found that the current FCC regulations [2] may be inadequate to protect TV viewers from the harmful effects of aggregate interference because the rules are made implicitly with only a single transmitter in mind. We suggested setting a maximum power density for secondaries rather than a per-device power limit to avoid this problem and verified its performance.

Given that enforcing a power density is necessary for TV receiver protection, we then explored how this solution can be employed to improve utility as well as freedom for secondary users. The current FCC rules implicitly favor certain applications. We offered a principled way to help mitigate the tensions between two types of users, rural and urban, using an approximately-optimal algorithm to choose a power density. The algorithm takes advantage of the logarithmic shape of the rate function. Unfortunately, this rule is not completely agnostic of usage case so there remain unresolved tensions.

The TV whitespaces represent a new paradigm for spectrum usage and as such are full of challenges and opportunities. The toolkit we built was designed to be flexible and modular enough to enable us to explore the TV whitespaces and address some of the many interesting questions that remain unanswered:

- What are the limits of our power density model?

- How should we certify devices for power-density compliance?
- When is spectrum sharing the optimal choice if a power-density scaling rule is present?
- How can we best navigate the tensions between short-range and long-range users?
- How can we further exploit the databases?

Appendix A

Methods

In order to understand the whitespaces on a nationwide scale, we developed a toolkit with the ability to let us explore using real-world data. We collected data from three major sources and created an interface for each. This library allowed us to build more complex models to answer some of the following questions:

- How many channels are available to secondaries?
- What data rates might be achievable in the TV whitespaces?
- How does the choice of communication range affect secondaries?
- Are whitespaces the right choice? How do they compare to traditional channel reassignment?
- Are TV receivers inside the protected region actually protected?
- How can we design the rules to be agnostic of secondary deployment?

This section covers the methodology behind the creation of the data presented in this thesis. Limitations of this data and our use of it are discussed in Appendix B.

It is worth noting that there were two major challenges in developing this project: speed and organization. There are many ways to answer these questions, but not all of them are fast in terms of execution time. It is important to find fast implementations in order to keep execution time down as much as possible. Having a rapid prototype process is key to developing ideas. Further, there is so much data generated in the process of making the figures featured herein–and even that is only a fraction of what is generated before the final version–so it is extremely important to keep everything organized so the results are accurately reported.

Further information as well as the toolkit itself can be found online at [6]. Both are updated as often as possible and inquiries are answered as quickly as possible.

A.1 Basics

The code to generate all data figures is written and executed using a basic installation of Matlab (no toolboxes).

To create map-level data, we use a discretized version of the continental United States. In particular, we discretize a Mercator projection of the US, as shown for the state of Wisconsin in Figure A.1. These points are sometimes referred to as "pixels." For the map size used, pixels vary in area from 210 km^2 to 290 km^2 (the Mercator projection is not an equal-area projection).

While this discretized version of the map gives a good high-level version of events, it fails to capture many local features. We will give one example of how to overcome this in Section A.6.

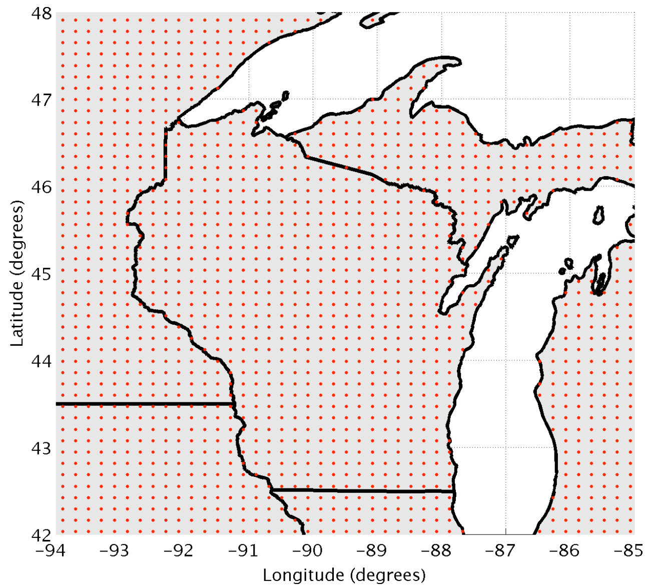

Figure A.1: Zoom-in of discretized map

A.2 External data

The simulations rely on three key pieces of external data, each of which will be discussed in turn:

- International Telecommunication Union (ITU) signal propagation model
- TV assignment data from the FCC's Consolidated Database System (CDBS)
- Population data from the US Census Bureau

Mishra, et al., were the first to develop a toolkit which incorporated these three elements [40, 41]. Our work builds on theirs and appears to have inspired similar work for Europe [53, 37]

A.2.1 ITU propagation model

The ITU propagation model uses transmitter height, power, and channel to determine the signal strength at any location[1]. This data is used for any plot which involves a protected radius, signal strength, or data rate. For example, all nationwide maps and CCDFs use this data. One can see instantiations of this propagation model in Figure 2.4 and online at [17].

The propagation model data was collected by Mubaraq Mishra from [10].

A.2.2 TV assignment data

TV assignment data is used to calculate pollution, exclusions, lost TV channels (jam motivation), and jam levels.

[1] Appendix B discusses why this is not strictly true.

The data used was downloaded from [3] on 16 September 2011. This website is a front-end for the FCC's CDBS, the resource used by TV Bands Database Administrators [18]. It includes geographic coordinates, transmit power, HAAT, channel, analog vs. digital, and much more for each TV assignment.

A.2.3 Population data

Population data is used to give per-person data rate estimates (e.g. maps) as well as to calculate per-person CCDFs.

Tract-level 2010 population data was obtained from the US Census Bureau [4]. The Census Bureau describes tracts as follows [5]:

> "Census tracts are small, relatively permanent statistical subdivisions of a county. Census tracts are delineated for most metropolitan areas (MA's) and other densely populated counties by local census statistical areas committees following Census Bureau guidelines (more than 3,000 census tracts have been established in 221 counties outside MA's). Six States (California, Connecticut, Delaware, Hawaii, New Jersey, and Rhode Island) and the District of Columbia are covered entirely by census tracts. Census tracts usually have between 2,500 and 8,000 persons and, when first delineated, are designed to be homogeneous with respect to population characteristics, economic status, and living conditions. Census tracts do not cross county boundaries. The spatial size of census tracts varies widely depending on the density of settlement."

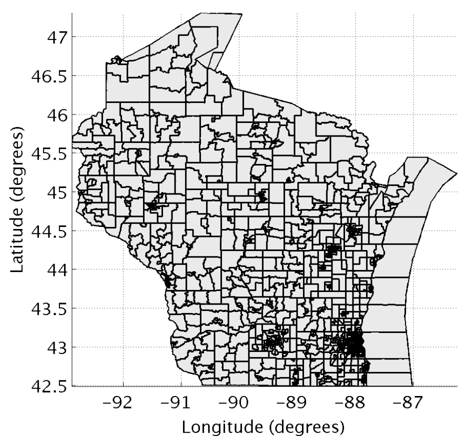

Figure A.2: Zoom-in of discretized map

Because of the way the census tracts are created, their size is directly correlated with their population density. For example, we can easily identify the urban areas of Madison and Milwaukee by the clusters of small tracts.

Also note that some census tracts extend beyond the land border of the state and into the water (e.g. Lake Michigan).

In order to calculate the discretized version of the population, we first calculate the population of the region represented by each pixel, shown in blue in Figure A.3. Each intersecting census tract contributes a percentage of its population proportional to the percentage of its area included in this pixel. For example, a tract which has 25% of its area inside of a particular pixel will add only 25% of its population to the running total for the pixel. Since census tracts are designed to be roughly homogeneous [5], this is a reasonable assumption. Population density is calculated by dividing a pixel's population by the area it represents.

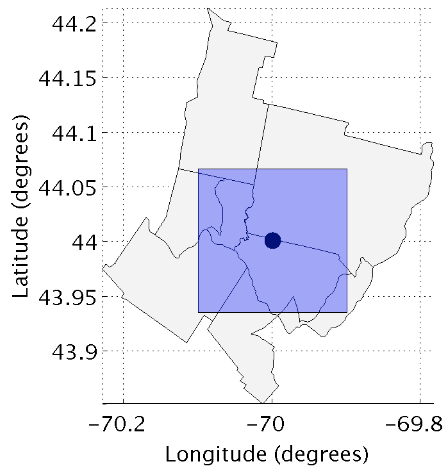

Figure A.3: Population computation illustration

Unfortunately, pixels are too large to sufficiently capture all of the variation in population. See Section B.2 for further discussion of this limitation.

A.3 Calculating single-link rates

Figure A.4 depicts the flow of information used to create a single-link rate map. Green blocks are input data, blue blocks represent intermediate pieces of data, and orange represents the final product.

First, we compute the noise map, signal map, and exclusions map separately:

- To create the noise maps, we use TV assignment data (location, transmit power, HAAT, and channel for each TV assignment) in conjunction with the propagation model to find the cumulative TV signal strength at each point on our map.
- We calculate the exclusions map, again using TV assignment data, propagation model, and the FCC's rules[2].
  - For digital TV assignments, we calculate using the F(50, 90) curves which means that the signal strength exceeds the threshold at 50% of locations 90% of the time.
  - For analog TV assignments, we calculate using the F(50, 50) curves.
- We create the signal map using the secondary's transmit power, HAAT, and channel to find its signal strength at its receiver. Note that in all but the $p = 2000$ case, the range is constant throughout the nation and thus the signal map is also constant. In the $p = 2000$ case, the variation in signal strength is due to the variation in ranges, which in turn are directly related to the local population density.

We divide the signal map by the noise map to yield the secondary's SNR map, then use Shannon's capacity formula to find the rate map.

Finally, we apply exclusions to the capacity to get the final rate map. Note that we can apply the exclusions last since pixels do not interact in this example (which is true for all but the power-density model).

[2] See Appendix B for caveats including HAAT calculation and propagation model for rule calculations.

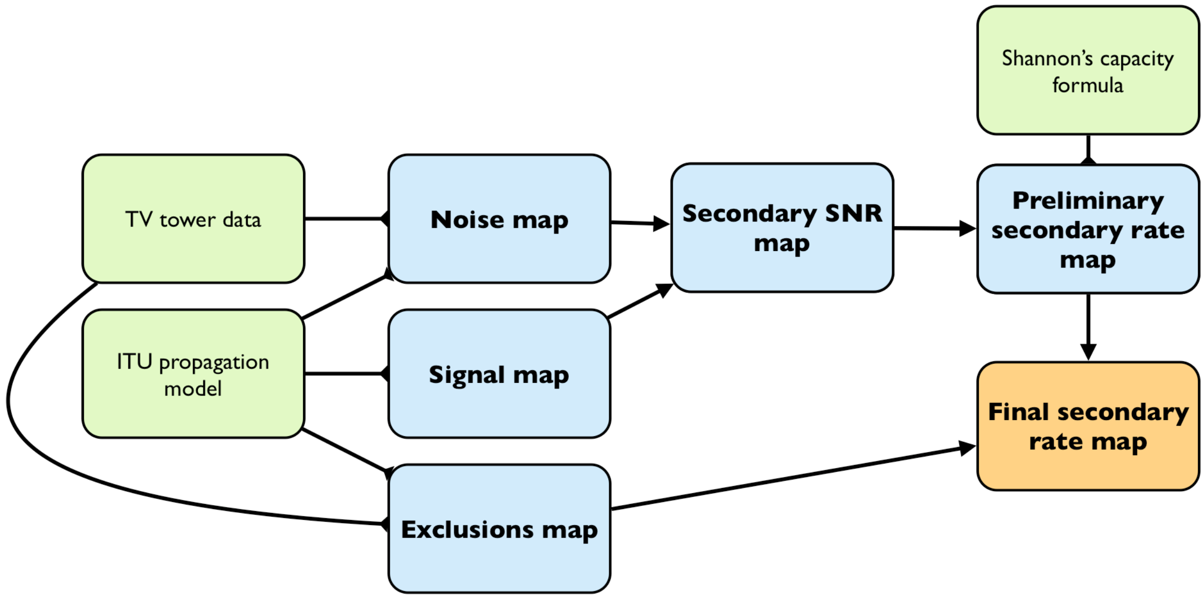

Figure A.4: Flow of information used to create a rate map for the single-link case.

A.4 Calculating a cumulative distribution function (CDF)

The cumulative distribution functions (CDFs) we calculate are weighted CDFs, typically weight by population but sometimes weighted by area. To calculate the CDF, we need two pieces of information: a weight map and a data map. The weight map assigns a weight to each corresponding point in the data map.

When the weight map is the population map, this can be viewed as giving each person a "vote." Once all of the votes are tallied, we create an empirical CDF of vote values. From this CDF we can easily calculate the empirical mean and median.

A.5 Computing MAC model rates

We use a medium access control (MAC) model to explore the effects of self-interference among secondaries. Our model allows secondaries to specify the size of the surrounding area which is free from other secondary transmissions, also called the *MAC exclusion area*. Outside of this exclusion area of radius r, we assume that the receiver is surrounded by neighbors as depicted in Figure A.5. In our MAC scheme, any transmitter closer than r km (the *MAC exclusion radius*) from our receiver of interest must halt transmissions[3]. Thus we can consider a *capacity per area*, defined as the rate that this receiver achieves divided by its footprint, $\pi \cdot r^2$. Finally, the value for r is chosen such that the capacity per area is maximized. Choosing a smaller r would reduce the size of the exclusion area but increase the interference levels, causing the capacity per area to shrink. A larger r would decrease the interference levels at the cost of enlarging the exclusion area, also causing the capacity per area to drop.

[3]One can imagine using some time-sharing protocol to ensure fairness among users, such as TDMA [30].

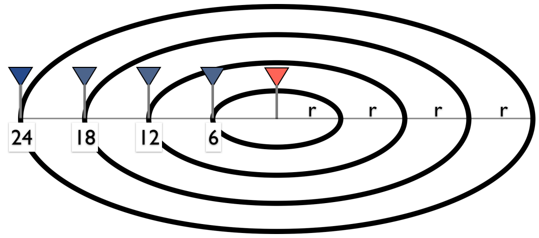

Figure A.5: Secondary locations in the MAC model: the receiver of interest (in red) is surrounded by 6 neighbors at distance r, 12 neighbors at distance $2r$, and so on.

A.6 Computing power density levels and rate maps

As is shown in Figure A.6, the steps can be broadly grouped into the categories "rule generation" and "rule evaluation." Because of this separation, we can calculate with one usage case in mind and evaluate with another.

We first calculate the allowed power density for each tower as if it were the sole TV tower[4]. We then conservatively combine these per-tower power densities by using the minimum allowed power at any location. Next, assign powers to secondary cells by multiplying the local power density by the cell's footprint. We then calculate the self-interference from nearby secondaries and use the final SNR to calculate the fair rate.

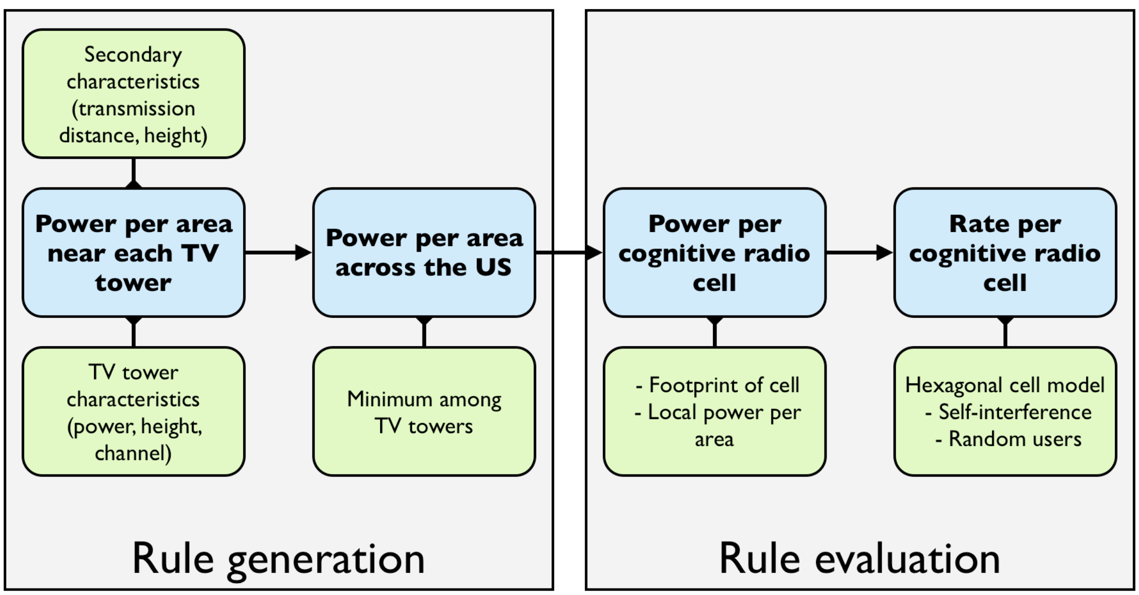

Figure A.6: Flow of information used to create power density maps and rate maps.

These steps are described below; however, many implementation details have been omitted in favor of imparting a high-level understanding. The interested reader will find these details online at [6].

Per-tower power calculations

The per-tower power calculations can be broken down into the following steps:

[4]We also neglect the effects of the borders. For example, even a TV tower in San Francisco is assumed to be completely surrounded by secondaries.

- Assume that each TV tower is surrounded by a sea of secondaries which extends “infinitely” far outward from r_p. Practically, we use 750 km[5] and approximate the sea as a series of secondaries located on rings centered around the TV tower. Symmetry implies that secondaries on the same ring are equivalent, thus drastically reducing the number of required computations by considering a “representative” secondary.
- The representative secondary from each ring calculates its “dream” power and rate. Recall that the “dream” power for ring n finds the safe uniform power density assuming that only secondaries at ring n or further may transmit.
- Secondaries whose “dream” rates are too low (below the threshold[6] β) are no longer allowed to transmit. This threshold effectively defines the value of $r_n - r_p$.
- We then find the maximum γ such that the total interference to a TV receiver at r_p is less than thermal noise[7].

Now we know the maximum safe power as a function of distance to r_p for towers individually. The next step synthesizes this data into nation-wide safe power levels.

Map-level power calculations

For each (channel, pixel) pair, we calculate the minimum power density allowed given the range from each of the TV transmitters on this channel. We call this the cochannel allowed power density for channel c, P_c.

We assume that TV receivers are able to attenuate adjacent-channel signals by 50 dB. This gives rise to a new adjacent-channel “exclusions” rule. Each (channel, pixel) pair is only allowed to use power density $\min\{P_c, P_{c+1} + 50\text{dB}, P_{c-1} + 50\text{dB}\}$.

For each location on our map we now have a set of (channel, safe power density) pairs. This completes the rule generation and our next step is to evaluate the rule using the data-rate metric.

Map-level rate calculations

As in the cellular model, the size of the cell depends on the local population density and the value of p. The total allowed transmit power is determined by multiplying the cell area by the local power density.

The calculations of desired signal strength and pollution level are identical to those in the cellular model. However, the self-interference calculations are more involved. We divide our interfering neighbors into two categories, distant and nearby, and add together the interference from each.

- *Distant transmitters*. Since the distance is so great, we can approximate the noise from far-away transmitters by calculating the interference from the center of their pixel to the center of our own. Since pathloss computations are time-consuming, this greatly reduces the run-time.
- *Nearby transmitters*. For nearby interferers the pixel-to-pixel approximation breaks down because the relative error is much higher. Instead, we use nearby power densities in our cellular model which is much more accurate at this smaller scale.

Now that we have the necessary pieces of information, we can calculate the SNR and therefore the rate of each secondary. As in the cellular model, we use the harmonic mean of all users in the cell to represent the fair rate per person of the cell. We now know the achievable data rate available to cellular-style deployments under the power-scaling rule.

[5] We found that results vary little with a “world” of this size.
[6] This is formally defined in Section 4.2.3.
[7] Under our assumptions, this is the maximum tolerable interference.

Appendix B

Limitations

Here we detail the assumptions and limitations of our toolkit.

B.1 Propagation model

- For simplicity and time, we completely ignore the effects of terrain and horizon. Mähönen, et al., showed that this will likely change local details but not nationwide trends [53].
- We use the ITU propagation model [10] rather than the Section 73.699 propagation model [17] for which the FCC regulations [2] are intended.
- The ITU model cannot calculate the pathloss for distances greater than 1000 km nor for heights below 10 meters (roughly rooftop-height).

B.2 Population

- We assume that the population of a pixel is roughly uniform. However, the population of the United States varies too quickly for this to be true. We show the magnitude of this discrepancy in Figure B.1 where we plot the CCDF by person of the population density itself. This shows that the statistics of our sampled population distribution do not match those of the true distribution[1]. Even with a 16-fold increase in sampling frequency we do not match the true distribution. It is infeasible to compute such high-resolution maps, so we leave it as future work to show the magnitude of this error after it propagates through our various models.
- Two major regions are missing population data. The first is in southern Nevada and the second is in southwest Arizona. These are shown as black regions on the population map in Figure 4.4. We believe that these are the locations of military bases.

B.3 FCC rule interpretations, towers

- The FCC regulations specify calculating the protected contour using eight radial arms extending from the TV transmitter. For each, the radial height above average terrain (HAAT) is calculated, rounded

[1] This was calculated using census-tract-level data which we assume to be the ground truth.

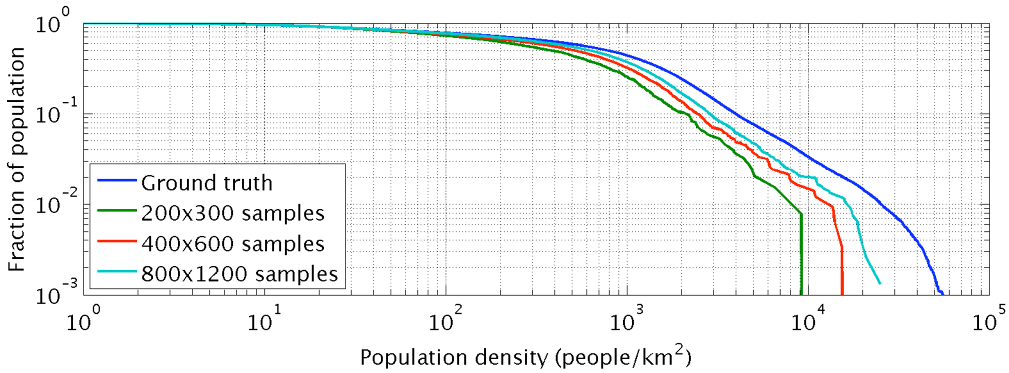

Figure B.1

up to 30 meters if necessary, and then used in the protected contour calculation. We use a single HAAT which we round to 10 meters (the smallest possible in the ITU model) if necessary. Because we use a single HAAT, our protected region is circular rather than irregular.

- We ignore all radio astronomy exclusions [2, §15.712(h)]. The locations are shown in Figure B.2. For visibility the circles shown have a radius of 24 km, ten times that of the actual exclusions. It is unclear how to account these excluded areas in our power-density model.

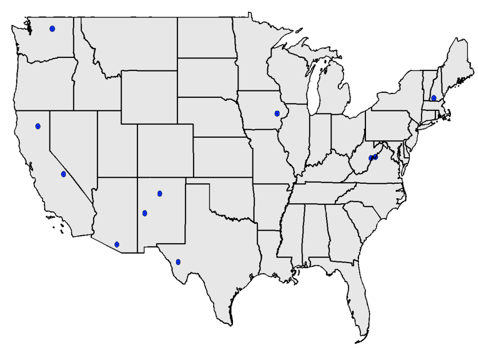

Figure B.2: Map of the United States showing the locations of the protected radio astronomy sites. For visibility, the protected regions are not shown to scale: the actual excluded regions are 1/100th the area shown. No secondaries may transmit in these regions.

- We ignore the PLMRS/CMRS exclusions in the 13 major metropolitan areas [2, §15.712(d)]. The number of reserved channels is shown in Figure B.3a. When we consider that some channels are already "lost" to TV tower exclusions, we see that the effect of including the PLMRS/CMRS exclusions is not so drastic; this effect is shown in Figure B.3b.
- We ignore border exclusions. The FCC rules state that foreign protected contours need only be protected in their country of origin [2, §15.712(d)].
- We assume reception threshold is 15dB, even for analog stations. In fact, except for calculating r_p, we assume all stations are digital.

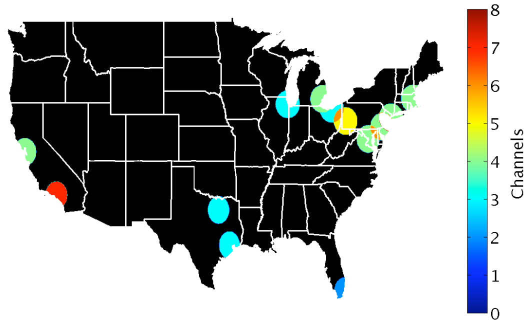

(a) Reserved PLMRS/CMRS channels in the 13 major metropolitan areas

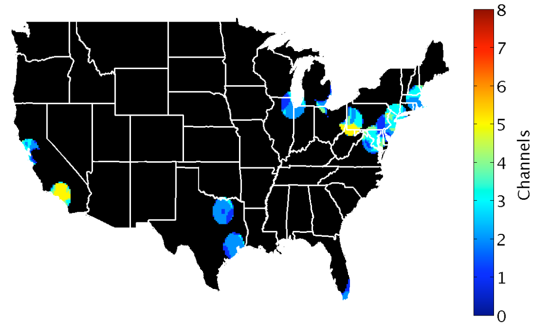

(b) Number of secondary channels excluded *only* because of PLMRS/CMRS exclusions

Figure B.3: Maps of the United States showing the effect of the PLMRS/CMRS metropolitan area exclusions.

B.4 Power density model

- The population density used for calculating the power density level near
- In order to calculate the "dream" rate and the "new" rate for each tower, we need to assume a transmission range. In the p-people-per-tower models, we assume that the local population density is the average of the population densities inside the TV station's protected contour. In fact, this is exactly the population that will not be served by the secondary devices for which we are calculating the power density. Furthermore, TV stations are located near population centers which means that the population is changing rapidly. We have preliminary results suggesting that this could cause improper power-density scaling but we have not yet remedied the problem.
- We assume that the TV decodability threshold is 15 dB.

Bibliography

[1] In the Matter of Unlicensed Operation in the TV Broadcast Bands: Second Report and Order and Memorandum Opinion and Order," Federal Communications Commision, Tech. Rep. 08-260, Nov. 2008. [Online]. Available: http://hraunfoss.fcc.gov/edocs public/attachmatch/ FCC-08-260A1.pdf.

[2] "Memorandam Opinion and Order on Reconstruction of the Seventh Report and Order and Eighth report and Order," Federal Communications Commision, Tech. Rep. 08-72, Mar. 2008. [Online]. Available: http://hraunfoss.fcc.gov/edocs public/attachmatch/FCC-08-72A1.pdf.

[3] See `transition.fcc.gov/mb/video/tvq.html`, in particular `http://transition.fcc.gov/fcc-bin/tvq?state=&call=&arn=&city=&chan=&cha2=69&serv=&type=3&facid=&list=4&dist=&dlat2=&mlat2=&slat2=&dlon2=&mlon2=&slon2=&size=9`.

[4] See `http://factfinder2.census.gov/faces/nav/jsf/pages/index.xhtml`.

[5] See `http://www.census.gov/geo/www/cen_tract.html`.

[6] See `http://inst.eecs.berkeley.edu/~harriska/CR/`.

[7] Barnouw, Erik: "A Tower in Babel", New York, Oxford University Press, 1966.

[8] U.S.C. Section 151.

[9] See `http://www.ntia.doc.gov/osmhome/allochrt.html`, `http://www.ntia.doc.gov/osmhome/allochrt.pdf`.

[10] "Method for point-to-area predictions for terrestrial services in the frequency range 30 MHz to 3 000 MHz," International Telecommunications Commission (ITU), RECOMMENDATION ITU-R P.1546-3, 2007.

[11] See `stakeholders.ofcom.org.uk/binaries/spectrum/spectrum-policy-area/spectrum-management/spectrum-usage-rights/sursguide.pdf`.

[12] See `http://www.comlaw.gov.au/Details/F2010L01990`.

[13] See `http://www.eecs.berkeley.edu/~smm/white_space_data_and_code.zip`.

[14] See `http://whitespaces.msresearch.us/WSWebGUI/whitespaces.aspx`.

[15] See `http://whitespaces.spectrumbridge.com/whitespaces/home.aspx`.

[16] See `http://hraunfoss.fcc.gov/edocs_public/attachmatch/DA-11-131A1.doc`.

[17] See `http://www.fcc.gov/encyclopedia/fm-and-tv-propagation-curves-graphs-sections-73333-73525-and-73`

[18] See `http://www.fcc.gov/encyclopedia/white-space-database-administration`.

[19] Technical and operational requirements for the possible operatin of cognitive radio systems in the white spaces of the frequency band 470-790 mhz, ecc report 159. `http://www.erodocdb.dk/docs/doc98/official/Pdf/ECCRep159.pdf`.

[20] L. Akter and B. Natarajan. Modeling fairness in resource allocation for secondary users in a competitive cognitive radio network. In *Wireless Telecommunications Symposium (WTS), 2010*, pages 1–6. IEEE, 2010.

[21] J. Bater, H.P. Tan, K.N. Brown, and L. Doyle. Modelling interference temperature constraints for spectrum access in cognitive radio networks. In *Communications, 2007. ICC'07. IEEE International Conference on*, pages 6493–6498. IEEE, 2007.

[22] D.B. Cabric. *Cognitive Radios: System Design Perspective*. PhD thesis, UNIVERSITY OF CALIFORNIA, 2007.

[23] M. Calabrese and B. Scott. Measuring tv 'white space' available for unlicensed wireless broadband. *New America Foundation and Free Press Analysis, January*, 2006.

[24] R.H. Coase. *The problem of social cost*. Wiley Online Library, 1960.

[25] M. Costa. Writing on dirty paper (corresp.). *Information Theory, IEEE Transactions on*, 29(3):439–441, 1983.

[26] A.S. De Vany, R.D. Eckert, C.J. Meyers, and D.J. O'Hara. Property system for market allocation of the electromagnetic spectrum: A legal-economic-engineering study, a. *Stan. L. Rev.*, 21:1499, 1968.

[27] J.P. de Vries and K.A. Sieh. The three ps: Increasing concurrent operation by unambiguously defining and delegating radio rights. In *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*, pages 56–70. IEEE, 2011.

[28] BP Freyens and M. Loney. Opportunities for white space usage in australia. In *Wireless Communication, Vehicular Technology, Information Theory and Aerospace & Electronic Systems Technology (Wireless VITAE), 2011 2nd International Conference on*, pages 1–5. IEEE, 2011.

[29] M. Gastpar. On capacity under received-signal constraints. In *In Proc 2004 Allerton Conference*. Citeseer, 2004.

[30] A. Goldsmith. *Wireless communications*. Cambridge Univ Pr, 2005.

[31] V. Goncalves and S. Pollin. The value of sensing for tv white spaces. In *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*, pages 231–241. IEEE, 2011.

[32] D. Gurney, G. Buchwald, L. Ecklund, S.L. Kuffner, and J. Grosspietsch. Geo-location database techniques for incumbent protection in the tv white space. In *New Frontiers in Dynamic Spectrum Access Networks, 2008. DySPAN 2008. 3rd IEEE Symposium on*, pages 1–9. IEEE, 2008.

[33] K. Harrison, S.M. Mishra, and A. Sahai. How much white-space capacity is there? In *New Frontiers in Dynamic Spectrum, 2010 IEEE Symposium on*, pages 1–10. IEEE, 2010.

[34] K. Harrison and A. Sahai. Potential collapse of whitespaces and the prospect for a universal power rule. In *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*, pages 316–327. IEEE, 2011.

[35] N. Hoven. On the feasibility of cognitive radio. *Master's thesis, University of California, Berkeley*, 2005.

[36] N. Hoven and A. Sahai. Power scaling for cognitive radio. In *Wireless networks, communications and mobile computing, 2005 International Conference on*, volume 1, pages 250–255. Ieee, 2005.

[37] R. Jantti, J. Kerttula, K. Koufos, and K. Ruttik. Aggregate interference with fcc and ecc white space usage rules: Case study in finland. In *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*, pages 599–602. IEEE, 2011.

[38] H.R. Karimi. Geolocation databases for white space devices in the uhf tv bands: Specification of maximum permitted emission levels. In *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*, pages 443–454. IEEE, 2011.

[39] M.J. Marcus. Unlicensed cognitive sharing of tv spectrum: The controversy at the federal communications commission. *Communications Magazine, IEEE*, 43(5):24–25, 2005.

[40] 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.

[41] S. Mishra and A. Sahai. How much white space has the fcc opened up. *IEEE Communication Letters*, 2010.

[42] R. Murty, R. Chandra, T. Moscibroda, and P. Bahl. Senseless: A database-driven white spaces network. In *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*, pages 10–21. IEEE, 2011.

[43] M. Nekovee. Quantifying the availability of tv white spaces for cognitive radio operation in the uk. In *Communications Workshops, 2009. ICC Workshops 2009. IEEE International Conference on*, pages 1–5. IEEE, 2009.

[44] J.M. Peha. Sharing spectrum through spectrum policy reform and cognitive radio. *Proceedings of the IEEE*, 97(4):708–719, 2009.

[45] C. Peng, H. Zheng, and B.Y. Zhao. Utilization and fairness in spectrum assignment for opportunistic spectrum access. *Mobile Networks and Applications*, 11(4):555–576, 2006.

[46] S. Pollin, B. Adams, and A. Bahai. Spatial reuse for practical scenarios: Iterative power adjustment from distributed contour estimation and propagation. In *Communications, 2008. ICC'08. IEEE International Conference on*, pages 2655–2661. IEEE, 2008.

[47] A. Sahai, K. Woyach, K. Harrison, H. Palaiyanur, and R. Tandra. Towards a "theory of spectrum zoning". In *Forty-seventh Allerton Conference on Communication, Control, and Computing*, 2009.

[48] C.E. Shannon and W. Weaver. *The mathematical theory of communication*, volume 19. University of Illinois Press Urbana, 1962.

[49] T.M. Taher, R.B. Bacchus, K.J. Zdunek, and D.A. Roberson. Long-term spectral occupancy findings in chicago. In *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*, pages 100–107. IEEE, 2011.

[50] R. Tandra, A. Sahai, and SM Mishra. What is a spectrum hole and what does it take to recognize one? *Proceedings of the IEEE*, 97(5):824–848, 2009.

[51] L. Tang, H. Wang, and Q. Chen. Power allocation with max-min fairness for cognitive radio networks. In *Mobile Congress (GMC), 2010 Global*, pages 1–5. IEEE, 2010.

[52] D. Tse and P. Viswanath. *Fundamentals of wireless communication*. Cambridge Univ Pr, 2005.

[53] J. van de Beek, J. Riihijarvi, A. Achtzehn, and P. Mahonen. Uhf white space in europe - a quantitative study into the potential of the 470–790 mhz band. In *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*, pages 1–9. IEEE, 2011.

[54] G. Vanwinckelen, M. Van Otterlo, K. Driessens, and S. Pollin. Reliable power control for secondary users based on distributed measurements. In *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*, pages 108–115. IEEE, 2011.

[55] M. Vu, N. Devroye, and V. Tarokh. On the primary exclusive region of cognitive networks. *IEEE Transactions on Wireless Communications*, 8(7), 2009.

[56] K. Woyach, P. Grover, and A. Sahai. Near vs far field: interference aggregation in tv whitespaces. 2011.

[57] K. Woyach and A. Sahai. Why the caged cognitive radio sings. In *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*, pages 431–442. IEEE, 2011.

[58] K.A. Woyach, A. Sahai, G. Atia, and V. Saligrama. Crime and punishment for cognitive radios. In *Communication, Control, and Computing, 2008 46th Annual Allerton Conference on*, pages 236–243. IEEE, 2008.

[59] H. Wu and Y. Pan. *Medium Access Control in Wireless Networks*. Nova Science Pub Inc, 2008.

[60] Y. Wu, E. Pliszka, B. Caron, P. Bouchard, and G. Chouinard. Comparison of terrestrial dtv transmission systems: the atsc 8-vsb, the dvb-t cofdm, and the isdb-t bst-ofdm. *Broadcasting, IEEE Transactions on*, 46(2):101–113, 2000.