How much white-space capacity is there?

Kate Harrison   Shridhar Mubaraq Mishra   Anant Sahai

harriska@eecs.berkeley.edu   smm@eecs.berkeley.edu   sahai@eecs.berkeley.edu

Dept. of Electrical Engineering and Computer Sciences, U C Berkeley

Abstract—**The November 2008 FCC ruling allowing access to the television white-spaces prompts a natural question. What is the magnitude and geographic distribution of the opportunity that has been opened up? This paper takes a semi-empirical perspective and uses the FCC's database of television transmitters, USA census data from 2000, and standard wireless propagation and information-theoretic capacity models to see the distribution of data-rates available on a per-person basis for wireless Internet access across the continental USA. To get a realistic evaluation of the potential public benefit, we need to examine more than just how many white-space channels have been made available. It is also important to consider the impact of wireless "pollution" from existing television stations, the self-interference among white-space devices themselves, the population distribution, and the expected transmission range of the white-space devices.**

The clear advantage of the white-space approach is revealed through a direct comparison of the Pareto frontier of the new white-space approach and that corresponding to the traditional approach of refarming bands between television and wireless data service. Finally, the critical importance of economic investment considerations is shown by considering the status of rural versus urban areas. Based on technical considerations alone, whether we consider long- or short-range white-space systems, people in rural areas would seem to be the main beneficiaries of white-space systems. A power-law distribution even appears that suggests that many rural customers could enjoy tremendous data-rates. However, the fundamental need to recover investments by wireless ISPs couples the range inversely to the population density. This clips the tail of the power-law and shows that urban and suburban areas can actually get significant benefit from the TV white-spaces.

Overall, the opportunity provided by TV white-spaces is shown to be potentially of the same order as the recent release of "beachfront" 700MHz spectrum for wireless data service.

I. INTRODUCTION

On November 14, 2008, the FCC released rules opening up the digital television bands to the operation of cognitive-radio devices [1]. In [2], we give an estimate for how much spectrum — measured in the number of channels — the FCC rules open up based on the 2000 USA Census and TV tower data extracted from the FCC database. Figure 1 shows the results in the form of a color-coded map of the continental United States showing how many MHz of white-space has been made available.

In [3], [4], we gave a systematic framework that helps in understanding the underlying policy dials: the FCC chooses an allowed transmit power for white-space devices and an "erosion margin" that determines how much extra interference to allow in the digital television bands, and this erosion margin determines which television receivers are deemed protected as well as how physically far away white-space devices must be from them (both on the channel itself and a separate distance
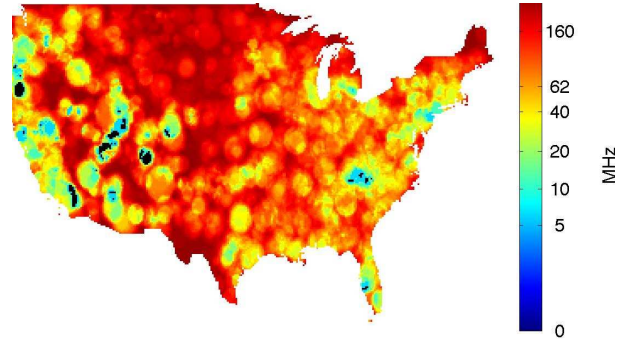

Fig. 1. A color-coded map of the continental USA with an estimate of the number of white-space channels allowed by the FCC's Nov 4th, 2008, ruling accounting for both co-channel and adjacent-channel protection. This map simply plots by latitude and longitude and does not use any other projection. The color legend is to the right and for comparison, the 62 MHz number is marked so that the white-space opportunity can be compared to the number of channels opened up in the 700MHz proceeding for wireless data providers.

for white-space devices operating on adjacent channels). In [3], the political tradeoff between the two candidate uses (broadcast television and white-space devices) was quantified in the native currency of politics: people. By looking at how many people on average gain access to white-space-channels as compared to how many people on average lose access to broadcast-television channels, we get a sense of the tradeoff between the two groups of users. In [3], it is shown that this political tradeoff is fundamentally better for white-space operation than it would be for the more traditional alternative of reassigning a channel from broadcast television over to hypothetical unlicensed wireless internet service providers.

However, a very important issue remained open in [2]–[4]: the relative importance of "pollution." While the above tradeoff emphasizes the issue of primary user protection, there is also the white-space device's perspective. As mentioned in [3], [4], the television white-spaces are not "white" in the sense of completely clean bands: they can have substantial "pollution" in them due to the presence of digital television signals. At first glance, this seems contradictory: after all, the white-spaces are locations where TV signals cannot be successfully received! However, there is no contradiction. As indicated in [5], the decodability threshold for digital television stations is about 15dB — that means that the TV signal is about **32 times stronger than noise** at the edge of where it can be decoded. Given the substantial height of television towers, the signal remains significantly stronger than noise for a long distance beyond that point. Thus, a straight comparison is not fair between the channels available to white-space devices and

those obtained by first kicking out televisions (as is the case in the 700MHz bands after the digital TV transition).

In [3], [4], this important issue was left open in the form of how much pollution a white-space device was willing to accept. There is more useful white-space available for devices that tolerate more pollution. However, this is hardly an acceptable point at which to leave the story. After all, spectrum is not itself a consumer good that can be directly enjoyed on its own terms by users. Instead, it is an input that is used by wireless systems to provide another intermediate commodity: data-rate delivered to the consumer. Diverse wireless-data applications in turn use the data-rate to enable delivery of desirable content actually enjoyed by the citizenry.

While an ideal tradeoff between TV and white-space devices would occur at the level of the desirability of the final content itself, there are many issues that make it hard to proceed in a definitive way along that path. After all, TV reception is multicast and presents only a finite set of choices to the consumer at any given time. Wirelessly delivered personalized content is drawn from a potentially much larger set of niches (compare what you can see by searching on YouTube versus flipping through the over-the-air channels right now). The complex realities of content distribution agreements and the illegal — while simultaneously ubiquitous — nature of much Internet-accessible content makes it hard to even begin crafting a meaningful comparison. Instead, we focus here on using the delivered data-rate itself to enable a decent evaluation of the value of the white-spaces.

It turns out that a variety of issues must be addressed to do such a comparison. And so, after reviewing some prior work on white-space evaluation in Section II, we show how to evaluate the white-spaces from a data-capacity point of view. The story is built up in stages using maps similar to Figure 1 to illustrate the magnitude and geographic distribution of the opportunity.

In Section III, we start with a pollution-only perspective that completely ignores the need to protect primary users, but does reveal the important role that the *range* of the wireless data system plays. The FCC's protection rules are added to the mix, and then the critical role of self-interference among white-space devices themselves is addressed to get a more realistic estimate of the data-rate available on a per square-kilometer basis. This captures the extreme personalization of Internet-style data-rate as contrasted with the mass-consumption delivery of television. Wireless-data users nearby must figuratively share the "tube" among themselves rather than consume the same content. For illustrative purposes, the map is then redrawn not in terms of data-rate, but in terms of the effective MHz in clean 700MHz that would be needed to achieve the same data-rate per area at the same range.

In Section IV, the key issue of the non-uniform distribution of population across the United States is introduced. The data-rate per area is normalized by the population density to give maps showing the per-person average data-rate available for different ranges. Curves are then shown that reveal the distribution of this data-rate over the population and these show a surprising new finding: that if the range is held constant, there exists a power-law distribution for the average data rate with a pretty heavy tail. The mean and median data rate differ by an order of magnitude.

Section V then switches perspective from the FCC's rules to the toy underlying policy tradeoff that is identified in [3]. The relative impacts of pollution, co-channel protection, and adjacent-channel protection are shown for short- and long-range wireless data service. The pollution-sensitivity of long-range communication is seen quite clearly. In addition, Pareto frontiers are illustrated comparing the average number of broadcast channels received by people to the wireless data-rate that can be delivered. The frontier expansion enabled by white-space operation as compared to traditional band reallocation is seen quite clearly.

The last part of this paper, Section VI, revisits the issue of range. When range is fixed, hyper-rural areas seem to get better per-person data rates. However, this misses the fact that communication range cannot be an exogenous variable when we take an economic perspective. A new model is used in which towers can only be built based on the number of customers available to amortize their operational costs. This results in higher tower-densities where the population density is higher and this has the effect of partially equalizing data-rates across the continental United States. The real losers are the very sparsely-populated areas: their economically viable ranges cannot technically support a high data rate.

Finally, a word on our methodology. Because of the intended audience of this paper, we do not dwell much on how the maps and plots were calculated. Standard models were used throughout and the actual Matlab source-code and raw data used is posted [6] to enable replication and follow-on work by other groups. Methodologically, there are a few limitations of this study that should be pointed out. Firstly, to calculate the available white space we assumed that all the licensed transmitters in the FCC high power DTV transmitters database [7] and the master low power transmitter database are all transmitting [8] and there are no other relevant transmissions[1]. As in [2], there is no counterpart in our study to some of the clauses from the FCC ruling: we neglected both wireless microphones and the more stringent emission requirements for the 602-620MHz bands (Section 15.709 [1]) while making this map. We also neglected the differences between the channel eligibility for fixed versus portable devices, and just assumed that all channels are available to fixed devices. We also neglected the locations of cable headends, fixed broadcast auxiliary service (BAS) links, and PLMRS/CMRS devices (Section 15.712 [1]). In addition, we assume that the ITU propagation models predict the reality on the ground to a fair degree [9]; this is particularly dubious in many areas due to the presence of mountain ranges, hills, etc.[2] Finally, we are overestimating the number of people served by broadcasters today by assuming that everyone in the noise-limited contour can and does receive a TV signal successfully.

---

[1]In particular, we are ignoring unlisted TV towers that might be across the border in Canada or Mexico.

[2]We also round HAATs for television towers up to 10m.

II. Prior work in estimating available white-space

There[3] has been prior work in estimating the amount of white-space available, but this has usually been done by lobbyists or by people who work for lobbyists. The problem is that both the language and methodology used by previous work does not properly distinguish between the pollution (which channels are attractive to use for us) and protection (which channels are safe to use without bothering others) viewpoints and this is the source of much confusion. There is even variation among those that focus on protection.

In [10], the author estimates that the average amount of white-space available per person is 214MHz. This is based on the FCC's estimate that the average American can receive 13.3 channels and there are 49 total DTV channels of 6MHz each. This line of reasoning tends to wildly overestimate white-space availability. For example, the FCC website [11] reveals that Berkeley, CA can receive 23 DTV channels. This would imply that the remaining 24 channels are available for white-space usage. However, only 5 channels are actually available for white-space use when the FCC's white-space rules are applied. This is because the FCC rules extend protection to adjacent channels and require a no-talk radius larger than the Grade-B protected contour [1]. Furthermore, low power TV stations and TV booster stations are ignored in [10], but these must also be protected.

New America Foundation also has another estimate of the amount of white-space available in major cities [12]. This study similarly overestimates the amount of white-space available (for example, they assume that 19 white-space channels will be available in San Francisco). Since the methodology for computing available white-space has not been described, we cannot explain the discrepancy between our paper and [12].

In [13], the authors claim to use the actual population data to quantify the amount of white-space available under different scenarios. The authors extract transmitter details from the FCC database (they did not use the High Power DTV transmitter list available from the FCC) as was done to prepare the plots in [14]. However, the authors assume that all locations beyond a station's protected contour can be used as white-space. The authors have estimated the amount of white-space under different scenarios. For scenario X (all DTV, Class A stations and TV translators; co-channel rules only) [13] estimates the median bandwidth to be ∼180MHz while our estimate in [2] of the median bandwidth per person is 126MHz. Similarly for scenario Z (all DTV, Class A stations and TV translators; co-channel and adjacent channel rules) [13] estimates the median bandwidth to be ∼78MHz while our estimate in [2] of the median bandwidth is ∼36MHz.

The major discrepancy can best be understood by a deeper inspection of the use of population density data in [13]. The authors estimate the size of each census block to be around 16 square miles. Each census block contains an average of 1300 people. For each transmitter the number of people in its protected region can be estimated by taking the number of census blocks that fit into its protected contour times 1300. While this may sound reasonable, the consequence is that the authors effectively assume a uniform population density across the United States! For such an artificial uniform population density, our estimate of the median bandwidth per person for scenario X is ∼186MHz while that for scenario Z, it is ∼102MHz. These numbers closely resemble the results in [13] (the slight discrepancy in these cases is due to the fact that [13] uses the FCC database which yields a much higher number of towers – 12339 versus 8071).

As far as we are aware, there has been no prior work on estimating the data-rates that might be available through white-spaces.

III. The basic story: capacity per unit area

A. A single link

Claude Shannon established a key relationship[4] in information-theory: $C = W \log_2(1 + SNR)$ where W is the bandwidth in MHz and the resulting data-rate is measured in Mbits/sec. For our purposes, W is 6MHz and the SNR term is the ratio of the received power from the desired signal to all the power in thermal noise and undesirable signals (what we call "pollution" here) received within this channel. The capacity adds across different channels and we use the propagation models specified in [1] here to calculate both the propagation for our desired signal (transmitting at 4W ERP in each 6MHz wide TV channel and from the maximum permitted height of 30m) as well as the pollution coming from the television stations registered with the FCC. The spillover from adjacent channels is assumed to be attenuated by 50dB within our white-space devices. The noise-figure is assumed to be perfect, but room-temperature thermal noise is still present as well.
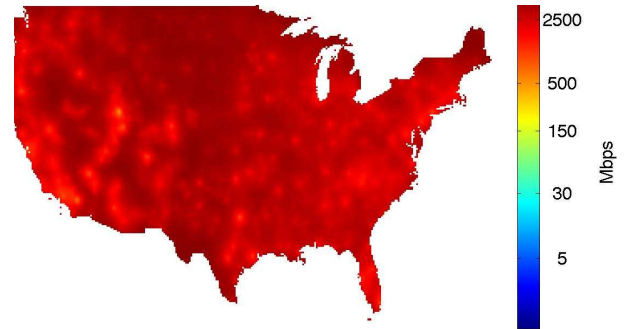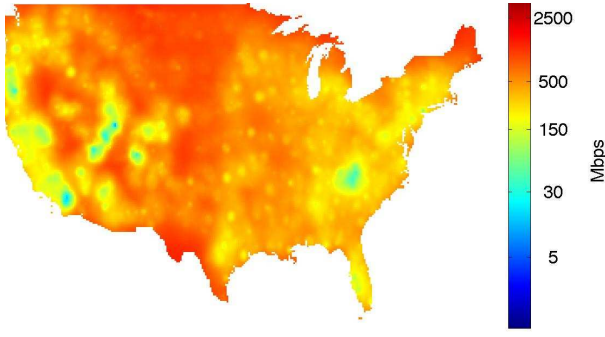

Fig. 2. A color-coded map of the continental USA with an estimate of the raw single-link capacity at a 1km range just treating the existing TV channels as pollution.

Figures 2 and 3 show the capacity distribution for a single isolated wireless link operating across the mainland at a link distance of 1km and 10km respectively. Notice just how much bigger the capacity is for shorter-range communication. This is due to the significantly stronger signal at a 1km range as compared to a 10km one.

Figures 4 and 5 only allow the use of those TV channels permitted under the white-space rules [1]: we are not permitted

---

[3]This section is largely copied from [2] to make this paper self-contained.

[4]Notice that here we are going to completely neglect the role of multipath fading and the possibility of using multiple-antennas to increase capacity and/or reduce interference.

Fig. 3. A color-coded map of the continental USA with an estimate of the raw single-link capacity at a 10km range just treating the existing TV channels as pollution.
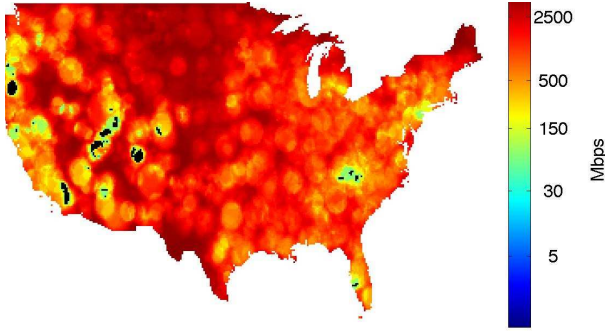
Fig. 4. A color-coded map of the continental USA with an estimate of the raw single-link capacity at a 1km range treating the existing TV channels as pollution and respecting the FCC white-space rules for protecting TV channels.

to use channels 3, 4, and 37 nor may we transmit within 14.4km of the protected contour of a co-channel or within 0.74km on adjacent channels. Notice that the FCC-rules do take a substantial bite out of the capacity, particularly for the short-range case. This is because at short-range, the received power from the white-space transmitter is high and so the strongly concave-∩ nature of the log function makes Shannon capacity relatively more sensitive to how many channels we have access to. By contrast, at long range, the received power

Fig. 5. A color-coded map of the continental USA with an estimate of the raw single-link capacity at a 10km range treating the existing TV channels as pollution and respecting the FCC white-space rules for protecting TV channels.

is low and so the log function is essentially linear around 1. Channels on which television can actually be received are then worth less than 10% of a clean channel in terms of capacity, and so their exclusion due to the need to protect primary users is not that painful. However, the need to protect adjacent channels is still painful since those would be attenuated by 50dB in terms of pollution.

B. A white-space network

The last section's ridiculously high capacities possible for a single link using white-spaces are misleading. This is because the economic value of the white-spaces is not in enabling one link, but in enabling coverage across the entire country. This means that the white-spaces need to be shared. In sharing, there are two effects. First, the signal from any given white-space tower is intended for one person at a time and hence the tower's capacity has to be divided by its footprint. While it is tempting to think that the white-space tower's footprint is just defined by its transmission range, this misses an important effect: the interference that a white-space receiver receives that is coming from *other* white-space users transmitting in the neighborhood. This is related to the important idea of *frequency reuse* in cellular systems [15]. Adjacent cells do not tend to use the same frequencies.
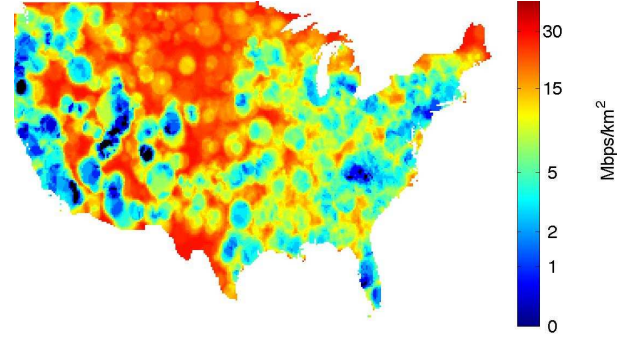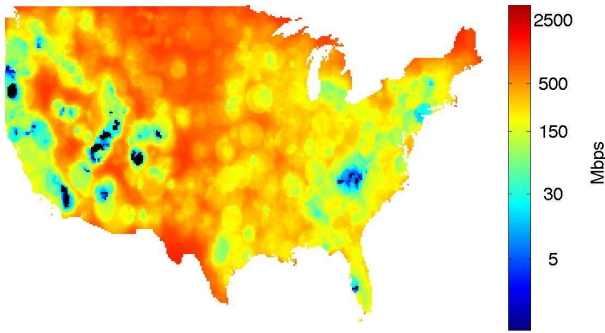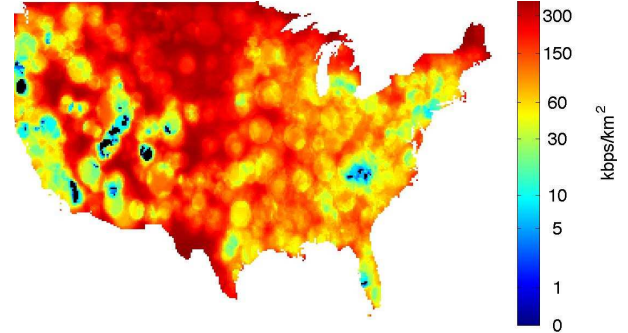

Fig. 6. A color-coded map of the continental USA with an estimate of the optimized capacity per square-kilometer assuming transmitters at a 1km range following FCC rules and optimizing the coexistence with neighboring white-space devices.

Fig. 7. A color-coded map of the continental USA with an estimate of the optimized capacity per square-kilometer assuming transmitters at a 10km range following FCC rules and optimizing the coexistence with neighboring white-space devices.

Formally, we define an exclusion-radius that tells other radios to keep out of this channel, and it is this exclusion-radius (not the range) that properly defines the footprint of the white-space tower from the perspective of resource sharing. This exclusion-radius can be optimized[5] to maximize the capacity per area. The results are illustrated in Figures 6 and 7. Notice the scales here: at 1km we are talking about rates in the MBits/sec per square kilometer and at 10km it is in the hundreds of kilobits/sec per square kilometer. The variation across locations is due to both the number of channels available *and* the differing amounts of pollution. The pollution level impacts the footprints: where there is a lot of pollution from TV signals, we do not mind having more nearby white-space devices either. This technical effect is, to our knowledge, new.

at 20W ERP on a clean channel to calculate the data-rate per square-kilometer that would be available. Figures 8 and 9 then show the effective number of such MHz that the white-spaces represent. Here, we see something that seems counterintuitive at first. Although TV channels are often touted as "beach-front property" in terms of their better propagation characteristics, the TV white-spaces turn out to be less valuable in these terms for longer-range because at that range, the pollution is also significant and turns out to dominate. However, the size of the opportunity is still quite significant.

IV. A HUMAN-CENTRIC PERSPECTIVE

In the end, white-space devices are going to be used by people. So, the population distribution needs to enter the picture. We used the Census data from the year 2000 that lists the population by zip code [16]. The zip code is specified as a polygon [17], and we assume the population is uniformly distributed[6] within that polygon. The white-space capacity per area can then be divided by the population density to get a long-term average capacity per person.

Fig. 8. A color-coded map of the continental USA with the effective number of MHz of spectrum opened up by the FCC white-space rules assuming transmitters at a 1km range.

Fig. 9. A color-coded map of the continental USA with the effective number of MHz of spectrum opened up by the FCC white-space rules assuming transmitters at a 10km range.

To compare the size of this opportunity to a known reference point, we take the recent 700MHz proceeding that released 62MHz of clean wireless data spectrum nationwide. Higher transmit powers are allowed and so we use a 40m high antenna
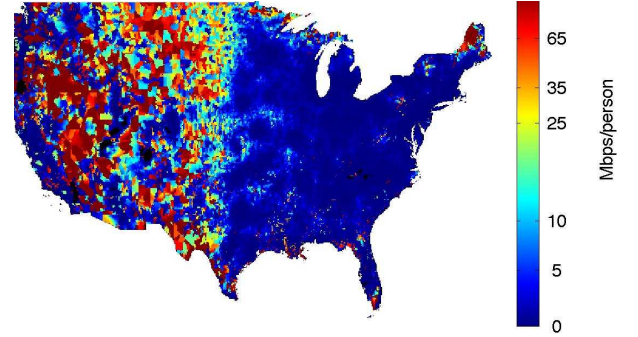

Fig. 10. A color-coded map of the continental USA with the capacity per person in the spectrum opened up by the FCC white-space rules assuming transmitters at a 1km range.
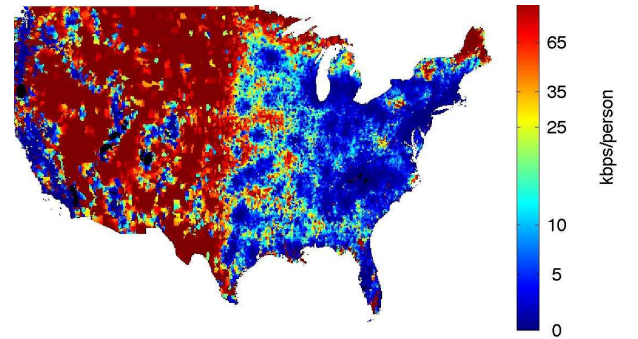

Fig. 11. A color-coded map of the continental USA with the capacity per person in the spectrum opened up by the FCC white-space rules assuming transmitters at a 10km range.

These per-person capacities are mapped in Figures 10, 11, and 12 for the white-spaces with a presumed 1km range, a

[5]A detail: in reality, interference does not just come from a single tower next door. It also comes from others at the same range. Furthermore, there are contributions from those that lie even further beyond, etc. Numerically, we optimize using a toy packing with 6 neighbors at a distance r, 12 further neighbors at a distance $2r$, 18 even further neighbors at a distance $3r$, and then 24 distant neighbors at a distance $4r$. Numerically, going beyond 3 rings makes very little difference because the signals have attenuated too far by then.

[6]So we are ignoring both the diurnal variation in population as many people commute to work and school as well as the finer structure of where residences are within each zip code.
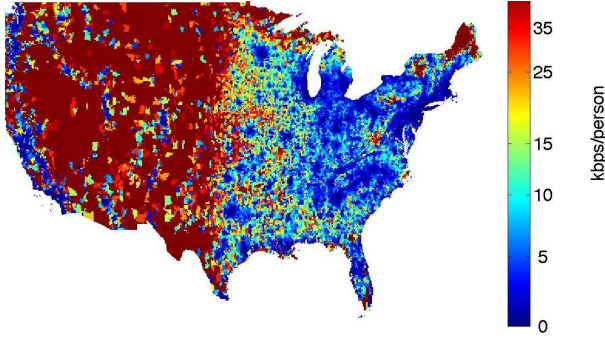
Fig. 12. A color-coded map of the continental USA with the capacity per person in the 62 MHz of 700MHz spectrum with transmitters at a 10km range.

Fig. 14. The distribution of data rate per person assuming 10km range to transmitters.

presumed 10km range, and for the clean 700MHz channels at a 10km range. The kbits/sec rates at the longer ranges can seem disappointing until we realize that these are long-term averages. If we assume that people only use the network to actively transport data for say 20 minutes in a day, then 10kbits/sec turns into a much more reasonable 720 kbits/sec while they are using it and about 3 gigabytes per month.

countrysides" where very few people live and so they get a very high data rate per person. The consequence of the power-law is that the mean and median are very different from each other — as we see quantitatively in the next section.

V. THE POLICY TRADEOFF

Fig. 13. The distribution of data rate per person assuming 1km range to transmitters.

Fig. 15. How the data-rate varies with the erosion margin that determines how much the TVs have to sacrifice for 1km-range wireless data service in the white-spaces.

The maps showing the qualitative behavior are complemented with probability distribution curves in Figures 13 and 14 that reveal the distribution across people and compare the white-spaces to the 700MHz bands. Notice the qualitative difference between the center and the tails. For most people, the 700MHz channels represent a larger opportunity than the white-spaces, although the difference is slimmer at a 1km range than at the 10km range. However, for the rare people who get high data rates, the white-spaces represent a bigger opportunity. This is because in these hyper-rural areas, there are also more channels available in general.

We also observe that there is a power-law that governs the per-person capacity. In a sense, this is the flip side of the well known power-law governing the population of cities. Except, this is a power-law that governs the *lack* of population in rural and wilderness areas. Just as there are far more mega-cities than the average would suggest, there appear to be "mega-
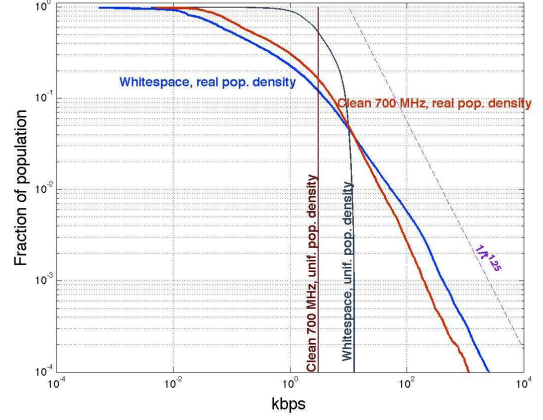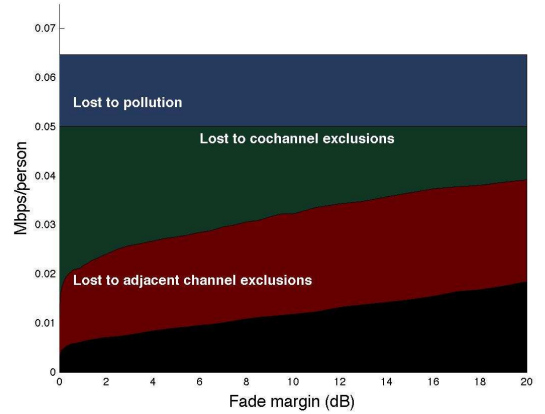
So far, the FCC rules have been taken as a single set of rules, not one possibility drawn from a family. As discussed in [3], [4], there is a natural way to parametrize the potential rules in terms of how much we allow white-space use to decrease the effective signal-reach of television transmitters. This is in terms of the erosion margin, and by varying it, we can see the range of possibilities in figures Figures 15 and 16. The vertical axis shows the median data-rate available on a per-person basis. The top represents what a clean channel would allow, and then we see the amount of data-rate lost to pollution, co-channel exclusions, and adjacent-channel exclusions before arriving at what median data rate we can deliver. Increasing the erosion margin can do nothing about pollution, but it does
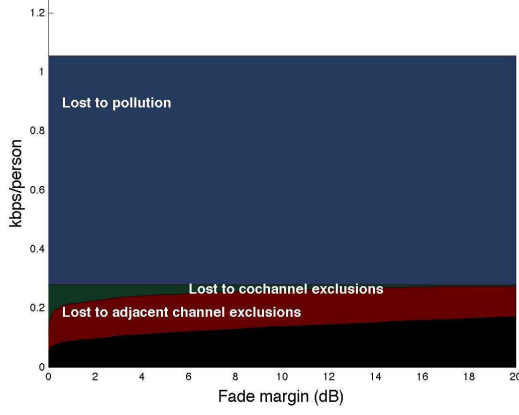
Fig. 16. How the data-rate varies with the erosion margin that determines how much the TVs have to sacrifice for 10km-range wireless data service in the white-spaces.
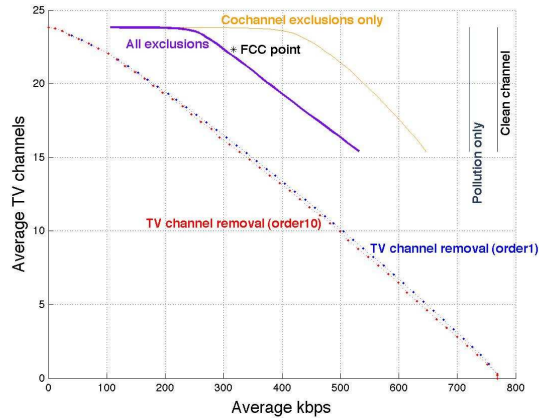

Fig. 18. The production-possibility frontiers for the tradeoff between TV viewers and *average* wireless data-rate for 10km-range wireless data service in the white-spaces.

diminish the losses due to the need to protect the primary user. Notice also the qualitative effect of the range: pollution is far more significant for long-range communication.
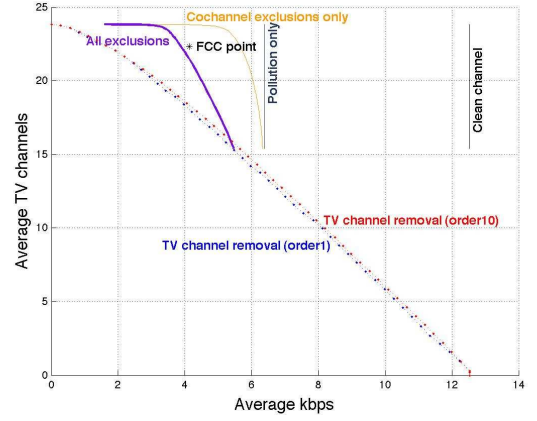

Fig. 17. The production-possibility frontiers for the tradeoff between TV viewers and *average* wireless data-rate for 1km-range wireless data service in the white-spaces.
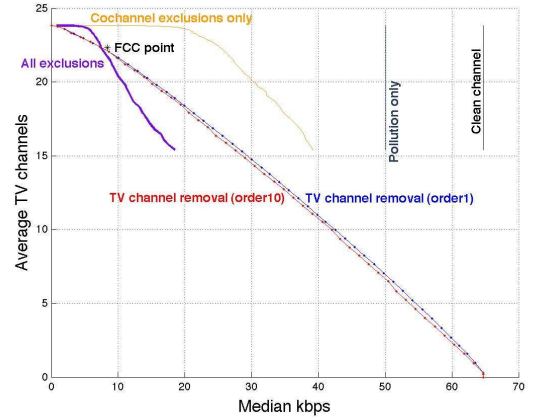

Fig. 19. The production-possibility frontiers for the tradeoff between TV viewers and *median* wireless data-rate for 1km-range wireless data service in the white-spaces.

The core tradeoff is better understood in terms of the average number of TV channels received and the data-rate received by the white-space device users. These are depicted in Figures 17 and 18 for the mean data-rate at 1km and 10km respectively, and in Figures 19 and 20 for the median data-rate. Notice that the mean data-rates are substantially higher (an order of magnitude) than the medians. This is a consequence of the underlying population power-law. In each case, the channels are removed in two orders: order1 and order10. Order1 is the sequence optimized[7] for the 1km case and similarly order10 is optimized for a 10km range.

The more significant observation is that the white-space approach can do better than the standard approach of taking
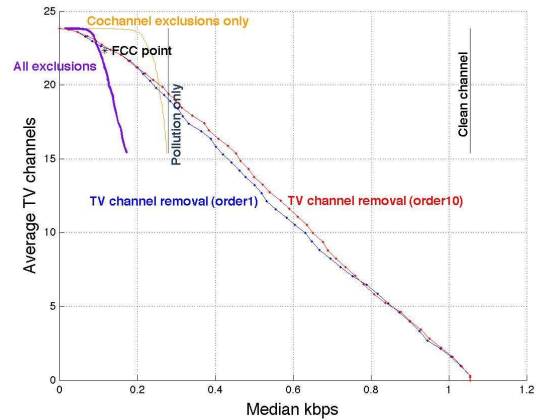

Fig. 20. The production-possibility frontiers for the tradeoff between TV viewers and *median* wireless data-rate for 10km-range wireless data service in the white-spaces.

[7]Channels are ranked according to their potential median data-rate (assuming thermal noise only) versus their number of TV viewers.

channels away from TV and giving them to wireless data providers — as was done in the 700MHz band. However, this only holds if we are unwilling to take away reliable reception for many TV channels. Once we are willing to lose a significant number of TV channels, it is worth just taking lightly-used channels away from TVs and reallocating them to wireless data.

Zoom-ups on the interesting corner of the Pareto frontier are shown in Figures 21 and 22. Here, the interesting thing to note is that the FCC's chosen point seems to reflect the interesting part of the white-space tradeoff. The fact that it lies off of our tradeoff curve is explained in [3], but the key reason is that here we are not assuming any directional antennas on the part of the primary TV receivers, while the FCC does. This allows the FCC to allow white-space device operation a bit closer to TVs than our model will allow because the TVs are capable of rejecting some of the white-space interference through their directional antennas. Notice that while the chosen FCC point is above the conventional refarming curves for the mean data rate at both ranges, it is located below the TV-channel-removal lines in Figures 20 and 22. So this issue of the mean versus median seems to be affecting our evaluation of the wisdom of the FCC's choice of tradeoff.
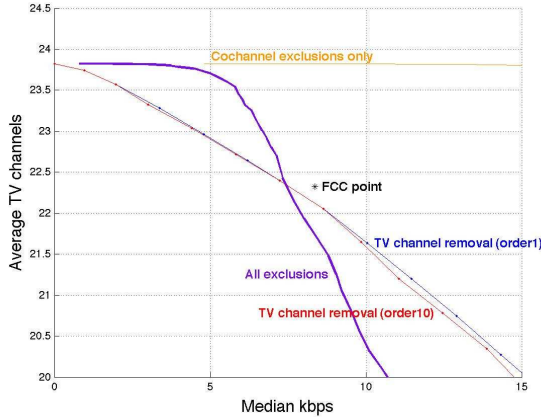

Fig. 21. Zoomed up production-possibility frontiers for the tradeoff between TV viewers and median wireless data-rate for 1km-range wireless data service in the white-spaces.

VI. ECONOMIC EFFECTS: INFRASTRUCTURE IS FOR PEOPLE

To get a better understanding, we must therefore decide whether the data-rate power-law is real or an artifact of our model. From a technical point of view, the deployment density of wireless data towers is an exogenous choice. However, from an economic point of view, it cannot be so. These towers are expensive and their costs must be shared over a base of customers. Where there are fewer people, we can only afford a few towers. Where there are more customers, we can place more towers. A uniform deployment density across a non-uniform population makes no economic sense.

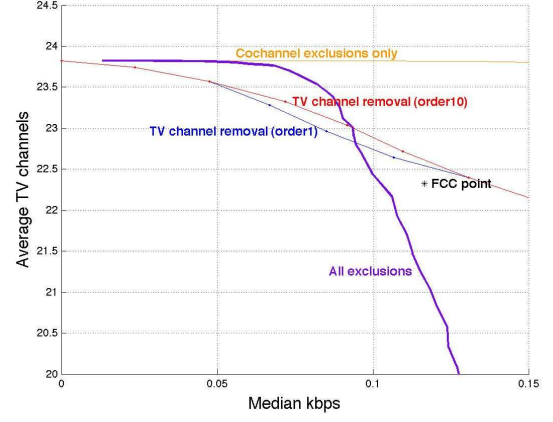For the purpose of illustration, we assume that it takes

Fig. 22. Zoomed up production-possibility frontiers for the tradeoff between TV viewers and median wireless data-rate for 10km-range wireless data service in the white-spaces.
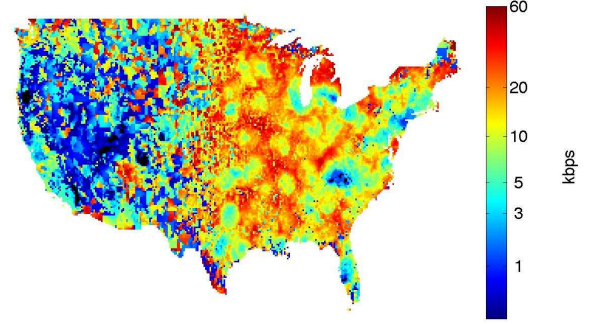

Fig. 23. A color-coded map showing the average data-rate available per person if the wireless range scales to preserve 2000 people per tower.

2000 people to support one tower[8] and cap the range to the nearest tower by 100km, even in the most remote regions[9]. Figure 23 shows the resulting capacity per person across the USA and the distribution is shown in Figure 24. Notice that the power-law is completely eliminated. The core reason for this is demonstrated in Figure 25 where we can see how even for clean channels, the capacity per person achieves a peak for each channel given a certain population density. Lower channels peak at lower densities, but the peaks are roughly in the same place.

In Figures 26 and 27, the tradeoff between TV viewers and data-rate is re-examined using this tower distribution model. TV channel removal has been optimized using the

---

[8]The guesstimate assumes that it costs $50K per year to build/operate a tower, families have 4 people in them, families are willing to pay an incremental $30 per month for white-space data service, and the wireless data providers want a healthy profit margin assuming 50% total market penetration.

[9]When using one tower per 2000 people, this cap affects approximately 10.4% of locations. However, at a range of 100km, the data-rate has already collapsed due to the exceedingly weak received signals. So this cap does not really matter.

same method as before. Notice that the mean and median data-rates are now of the same order of magnitude and agree qualitatively. We see in Figure 28 that in many areas this corresponds to roughly 62 MHz of spectrum, the same amount released in the recent 700MHz proceeding.
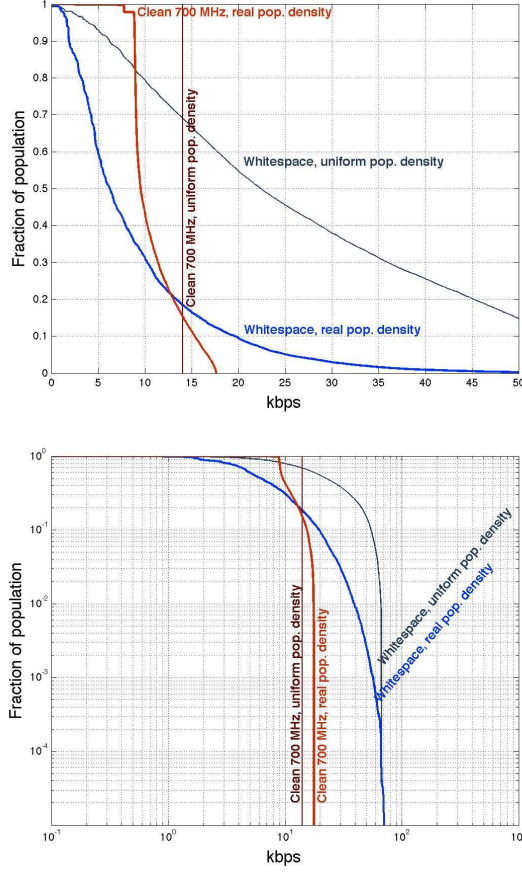

Fig. 24. The data-rate distribution available per person if the wireless range scales to preserve 2000 people per tower. The top is on a linear scale while the bottom is logarithmic.
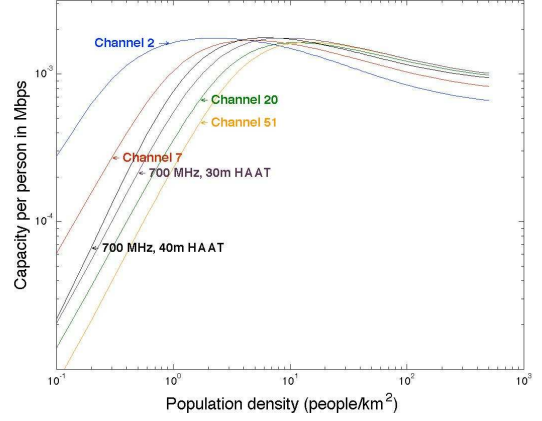

Fig. 25. For a clean channel, how the capacity per person varies with population density if we want to keep 2000 people per tower and adjust the range accordingly. Notice how the 700MHz number varies with antenna height.
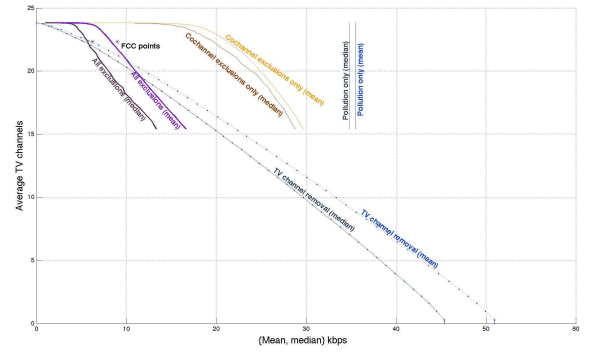

Fig. 26. The production-possibility frontiers for the tradeoff between TV viewers and wireless data-rate if wireless range scales to preserve 2000 people per tower.
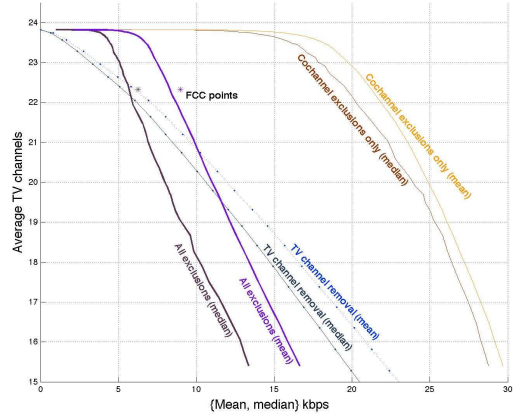

Fig. 27. Zoomed-up production-possibility frontiers for the tradeoff between TV viewers and wireless data-rate if wireless range scales to preserve 2000 people per tower.

REFERENCES

[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] S. M. Mishra and A. Sahai, "How much white space has the fcc opened up?" *To appear in IEEE Communication Letters*, 2010.

[3] ——, "Pollution vs protection in determining spectrum whitespaces: a semi-empirical view," *Submitted to IEEE Transactions on Wireless Communication*, 2010.

[4] ——, "How much white space is there?" Department of Electrical Engineering and Computer Science, University of California Berkeley, Tech. Rep. EECS-2009-3, Jan. 2009. [Online]. Available: http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-3.html

[5] 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," *IEEE Trans. Broadcast.*, vol. 46, no. 2, pp. 101–113, Jun. 2000.

[6] K. Harrison and S. M. Mishra, "White space code and data, version 0.2." [Online]. Available: http://www.eecs.berkeley.edu/~sahai/new_white_space_data_and_code.zip
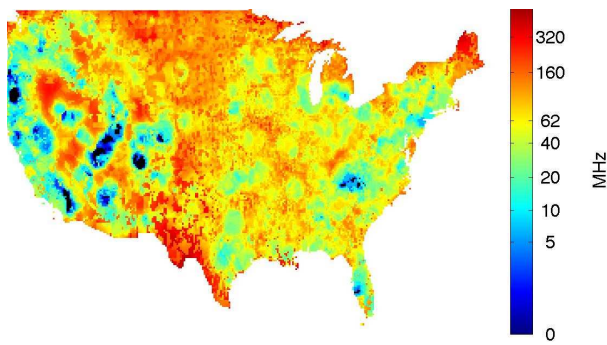
Fig. 28. A color-coded map of the continental USA with the effective number of MHz of spectrum available if the wireless range scales to preserve 2000 people per tower.

[7] "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

[8] "List of All Class A, LPTV, and TV Translator Stations," Federal Communications Commision, Tech. Rep., 2008. [Online]. Available: http://www.dtv.gov/MasterLowPowerList.xls

[9] "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.

[10] J. Snider, "The Art of Spectrum Lobbying: America's $480 Billion Spectrum Giveaway, How it Happened, and How to Prevent it From Recurring," New America Foundation, Tech. Rep., Aug. 2007. [Online]. Available: http://www.newamerica.net/publications/policy/art_spectrum_lobbying

[11] F. Communication Commission, "DTV Reception Maps." [Online]. Available: http://www.fcc.gov/mb/engineering/maps/

[12] B. Scott and M. Calabrese, "Measuring the TV 'White Space' Available for Unlicensed Wireless Broadband," New America Foundation, Tech. Rep., Jan. 2006.

[13] C. Jackson, D. Robyn, and C. Bazelon, "Comments of Charles L. Jackson, Dorothy Robyn and Coleman Bazelon," The Brattle Group, Tech. Rep., Jun. 2008. [Online]. Available: http://fjallfoss.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6520031074

[14] A. Sahai, S. M. Mishra, R. Tandra, and K. A. Woyach, "DSP Applications: Cognitive radios for spectrum sharing," *IEEE Signal Processing Magazine*, Jan. 2009.

[15] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*, 1st ed. Cambridge, United Kingdom: Cambridge University Press, 2005.

[16] United States Census, "USA Census 2000 Gazetteer files." [Online]. Available: http://www.census.gov/geo/www/gazetteer/places2k.html

[17] ——, "US Census Cartographic Boundary files." [Online]. Available: http://www.census.gov/geo/www/cob/st2000.html#ascii