

Spectrum Sharing in the Gray Space

Jon Peha¹

Carnegie Mellon University

Abstract

There has been considerable effort to let more wireless devices operate in “white space” spectrum, i.e. within frequency bands and geographic areas where no wireless devices are active. Making white space available is certainly useful, but there are other sharing opportunities as well, some of which have been obscured by dangerous misconceptions about the concept of “unused spectrum.” This paper discusses allowing more devices to operate safely in “gray space” spectrum, i.e. spectrum that is actively being used in that transmissions are underway – something many economic models assume is impossible. The paper focuses on primary-secondary sharing, so devices gaining access to spectrum operate on a secondary basis in a way that never causes harmful interference to primary systems. Examples of primary-secondary gray-space sharing mechanisms are described in which devices are allowed to share spectrum with broadcasting, radar, and cellular systems. Quantitative analysis shows that it is technically possible to support significant communications among secondary devices in spectrum that is already heavily used by cellular or radar. However, gray-space sharing generally causes primary and secondary systems to be more technically interdependent than white-space sharing, so different policy and governance structures are needed. Secondary market rules can support gray-space sharing in cases where there is a single primary spectrum user, such as a cellular carrier. In cases where technology is static, the regulator may be able to control access for secondary devices. However, in cases with multiple primary users and multiple secondary users of spectrum, as might be seen in bands with radar for example, a new kind of governance body will be needed to facilitate spectrum sharing.

Keywords: spectrum, gray space, white space, primary-secondary sharing, radar, cellular

¹ www.ece.cmu.edu/~peha/bio.html

Section 1: Introduction

There is growing pressure to support more and more communications within those spectrum bands that are most conducive to operating a cost-effective wireless system. Many people assume that the only way to meet this need is to find an “unused” block of spectrum, sometimes called “white space,” and make that block of spectrum available for use, where a block includes all the spectrum within a given frequency range and a given geographic area. This paper shows some fallacies and limitations of this white space approach, describes example mechanisms for “gray space” sharing in which devices are given access to spectrum that is already in use, quantitatively demonstrates the value of such schemes, and identifies policy and governance reform that could facilitate use of some of these gray-space sharing mechanisms. In particular, this paper focuses on primary-secondary (Peha 2009a) gray-space sharing, in which systems either have primary or secondary rights, and those with secondary rights are prohibited from causing harmful interference to those with primary rights.

This section begins with the murky and often counterproductive concept of “unused” spectrum. This term seems to imply that only one signal can *use* a block of spectrum at a time, but in reality, signals can easily coexist in space and time; signals are electromagnetic waves that simply pass through each other. This causes confusion among spectrum policymakers. For example, as part of a “spectrum inventory,” some people are calling for clear black-and-white maps that identify the locations where spectrum is “used” and locations where it is not in a given band. Proponents may support their call for such maps by citing measurements (Federal Communications Commission [FCC] 2002; New America Foundation & Shared Spectrum Company 2003; Shared Spectrum Company 2004) which they believe show that a small fraction of spectrum is “used.” In so doing, they are using contradictory definitions of unused spectrum. There are (at least) three distinctly different definitions. If one is trying to maximize spectrum sharing, the definition of unused spectrum that matters is spectrum that can support more use, i.e. additional devices can be deployed without causing or experiencing harmful interference. However, this definition is *highly* dependent on the specific type of device that would be deployed, so any map or statistic that is not specific to a device type is simply wrong. (A spectrum inventory could still be enormously useful (Peha 2009b), but it should be far more detailed than simply identifying spectrum as used or unused as some might wish.) Often, when people say there is unused spectrum, they mean there is a location and frequency at which no transmissions can be detected. This is of course how devices with spectrum sensors view the world, and it is what measurement studies actually indicate, but this is quite different from the first definition. As the radio astronomy community often points out (National Research Council 2010), one might detect no transmissions from a location very near important receivers, but allowing a transmitter to share spectrum there would cause serious harmful interference. Conversely, one might detect numerous transmissions, none of which would be harmed by the deployment of additional wireless devices. Indeed, some of what is detected is not a transmission at all, e.g. the output of a microwave oven. Moreover, any definition of unused spectrum that is based on whether transmissions can be measured must depend on how one does the measurements, including the threshold at which one defines a channel as busy, the size of that channel, and whether the characteristics of the signal to be detected are known. A third and probably the most meaningful definition of unused spectrum would be a geographic region and frequency range in which

neither transmitters nor receivers are operating above a certain threshold. In practice, when one says that devices can operate in this “white space” without causing harmful interference, this is typically what is meant. It is a conservative definition, meaning that allowing sharing in this kind of white space can be useful and effective, but there are other areas referred to in this paper as “gray space” in which devices are present, and their transmissions could be detected, but additional devices could still be deployed without causing harmful interference.

While the idea of allowing different systems to coexist in the same spectrum is not new, at least in the United States the current quest to free up spectrum for new purposes is largely focusing on two ideas centered around unused spectrum. Both approaches are worth pursuing, but should not subsume the entire agenda. One is the approach adopted in Nov. 2008 in the U.S. Federal Communication Commission’s (FCC) TV white space proceedings (FCC 2008). Under these rules, unlicensed devices are allowed to operate within frequencies allocated to television and within a geographic region that is far from any TV coverage area in this frequency range, as determined by a database that identifies entirely “unused” spectrum by location. Moreover, in the 2010 U.S. National Broadband Plan, the FCC further suggested that this particular approach should also be considered in bands other than those allocated to television (FCC 2010) as a means of relieving the shortage of available spectrum. For the proposed policies to work, conservative assumptions must be built. This database should prevent secondary devices from operating in buffer zones around the predicted TV coverage areas, in part to account for the fact that predictions are never exact. The database also does not allow devices to take advantage of changes in usage over time unless they occur very slowly and predictably, because doing otherwise would require unlicensed devices to check the database frequently. Thus, there are clearly additional spectrum-sharing opportunities.

A second approach that was explicitly highlighted in the U.S. National Broadband Plan (FCC 2010) uses secondary market rules (FCC 2003, 2004), which give a license-holder the right to lease a block of spectrum to others for short or long periods, which in this ruling means under or over one year. It is easy to imagine such a block of spectrum defined in geographic area and frequency range, just as many area-based licensed are defined today. This approach makes sense when the granularity of the blocks is large. For example, a license-holder has the spectrum to provide cellular service to a major city, but has no plans to use that spectrum, so it leases all rights to use of spectrum in and around that city to another company. There are also ways to make this spectrum available other than secondary markets, e.g. the regulator could allow the use of a licensed block of spectrum without the permission of the license-holder if and only if that block is found to be completely unused. This is the equivalent of allowing squatters in an empty property, but allowing the owner to reclaim the property at any time. However, this concept of an idle block of spectrum loses its meaning as the granularity becomes smaller, e.g. for cellular if the spectrum block covers a geographic area that is not significantly larger than the size of a cell. For example, if Device A wants to transmit to Device B, it does not need spectrum throughout a geographic area of any size; it only needs to know that the interference will be sufficiently low at Device B’s specific location, and that no receiver near Device A will experience harmful interference. No other location matters. Thus, the idea that each system needs its own block of spectrum with a defined frequency range and geographic boundaries, as assumed in the seminal paper

by De Vany et al (De Vany, Eckert, Meyers, O'Hara & Scott 1969) and many others since, is wrong in such cases. While it is still possible to support spectrum-sharing models over smaller areas and smaller time periods using secondary market rules (FCC 2004), it is necessary to look beyond the simple metaphor of leasing entirely unused blocks of spectrum.

There is certainly great value in putting spectrum that is truly unused to work, as both the approaches above do, but there are many ways to share spectrum and thus many ways to support more communications within a given block of spectrum (Peha 2009a). For example, reconsider this case of television. One gray-space approach is to allow secondary devices to operate within spectrum actively used for TV broadcasts, but only in those locations where the television signal is sufficiently strong that interference will not degrade TV quality. Thus, secondary devices might transmit if they are at outdoor locations reasonably close to the TV tower, even though TVs are active there as well. (Secondary devices could not transmit near the edge of TV coverage where the TV signal is weak.) Alternatively, secondary devices could operate *opportunistically* based on what they learn about current conditions (Peha 2009a). To do so, the secondary devices would need some means of determining whether there are TV receivers nearby that are on and tuned to that particular TV channel. If not, operating in this used spectrum would cause no harmful interference.

Section 2 discusses two forms of gray space sharing, and demonstrates that it is sometimes technically possible to support significantly more communications by sharing “used” spectrum. There is an endless variety of technical approaches to sharing and types of systems that could be involved. Section 2 considers two very different types of primary spectrum users as examples: cellular and fixed radar. A single cellular system typically blankets an entire region with coverage, and is given exclusive rights in the form of an area-based license to do so. In contrast, fixed radar facilities operate only in specific locations, and the radar systems in a given band may belong to different operators, as is also the case with television. As discussed in Section 3, policies for spectrum licensing and spectrum sharing should reflect these differences. Section 3 discusses policies that would facilitate gray space sharing, some of which are in place today in the U.S., and some are not. Conclusions are summarized in Section 4.

Section 2: Examples of Gray-Space Sharing Technology

This section presents two illustrative examples of primary-secondary spectrum sharing, in which primary and secondary systems operate in the same frequency band, the same geographic region, and at the same time, but the secondary user(s) are prohibited from causing harmful interference to the primary user(s) of spectrum. The primary system is cellular in Section 2.1, and fixed rotating radar in Section 2.2.

As this section will show, spectrally efficient gray-space sharing is possible, but it may take a different form from emerging white space sharing, and this has policy implications that will be discussed in Section 3. In general, primary-secondary sharing can be based on *cooperation*, which means the primary and secondary systems explicitly coordinate, or based on *coexistence*, which means they do not (Peha 2009a). With cooperative sharing models, both primary and secondary systems must be co-

designed to the extent that they support a common protocol, and primary systems typically have some direct control. With coexistent sharing models, it is possible to build secondary systems regardless of the types of primary systems with which they share spectrum, particularly in white space since by definition those primary systems will never be nearby when a secondary device is active. However, with the more spectrally efficient gray-space sharing, it is more likely that secondary systems must be designed to share spectrum with specific types of primary systems, and this is reflected in the examples below.

Section 2.1: Cellular

At least in densely populated areas, cellular carriers have strong incentive to make efficient use of their spectrum. Thus, they can be used to illustrate how heavily used spectrum can still be used further. Even spectrum that is 100% utilized from the perspective of the cellular carrier, which means it is not possible to support more cellular customers without building additional infrastructure, can be used for other purposes without degrading the quality of cellular service.

There are a number of approaches to support primary-secondary sharing with cellular (Bakr, Johnson, Wild & Ramchandran 2008; Lee, Han, Hwang & Choi 2009; Panichpapiboon & Peha 2008; Peha & Panichpapiboon 2004; Saruthirathanaworakun & Peha 2010). One possibility is to take advantage of the geolocation of primary and secondary devices (Panichpapiboon & Peha 2008; Peha & Panichpapiboon 2004). From the perspective of a cellular system, a transmission from a cellular tower to a cellular handset consumes some of the limited downstream capacity throughout a cell (or throughout a sector if that cell is sectorized), regardless of precisely where in the cell that handset is. However, if a secondary device were present in the cell, it would only harm the transmission to that handset if the secondary device were close to the handset. Thus, depending on the location of the active handsets in a cell, it can be possible to allow a secondary device to transmit from within a heavily used cell without causing harmful interference. Moreover, it is easy for a cellular carrier to move handsets from one frequency band to another within the carrier's assigned spectrum. Using these observations, a cellular carrier could offer a new service, whereby secondary devices request the right to transmit from a given location, and each request is either accepted and assigned a given frequency range or rejected, depending on whether the carrier can place all handsets near this secondary device in other spectrum bands. Note that the secondary device would then operate in cellular spectrum, but would not use cellular infrastructure in any way. Thus, this is cooperative sharing (Peha 2009a), with the cellular carrier in control. In this way, it is possible to support a large number of secondary devices in each cell with no degradation to cellular service, even when the cellular system is 100% utilized, provided that it is possible to preempt secondary transmissions when necessary if a cellular handset near the secondary device initiates a new call. Indeed, even if secondary devices are guaranteed the right to transmit without ever being preempted, which could make this arrangement far more attractive to secondary users than unlicensed spectrum, it is possible to allow dozens of 100 mW secondary devices per cell to transmit simultaneously and for extended periods in 10 MHz of cellular spectrum with only a modest impact on total cellular capacity (Panichpapiboon & Peha 2008; Peha & Panichpapiboon 2004).

Another option is to allow secondary devices to transmit opportunistically in spectrum used for upstream cellular communications, but only at times when the signal to interference + noise ratio (SINR) at the cellular basestation is sufficiently high that additional interference is tolerable (Saruthirathanaworakun & Peha 2010). In a cooperative scheme, the basestation could periodically announce the extent to which it can tolerate additional interference to any secondary devices in the cell. In a coexistent scheme, secondary devices could estimate the SINR at the basestation by monitoring downstream transmissions, and taking advantage of the high correlation between the rates of downstream and upstream transmissions (Saruthirathanaworakun & Peha 2010). Even when there is no cooperation between primary and secondary systems, designing a secondary system that can avoid harmful interference to cellular requires knowledge of the technology used by the primary system, including parameters like maximum transmit power, and maximum tolerable SINR. Thus, the technologies of primary and secondary systems are more intertwined than can be seen with TV white space devices, for example, for which design does not depend on how television transmitters or receivers are designed.

The benefits of this form of sharing with cellular have been quantified in the technical literature (Saruthirathanaworakun & Peha 2010). In one example scenario, the primary spectrum user is a CDMA-based cellular system operating at 880 MHz. Cell radius is 8 km, and cellular traffic follows a standard Erlang model with a tolerable call blocking probability of 2%. The secondary spectrum user in a given cell is a 400-meter point-to-point link that is 6.8 km from the tower. Fig. 1 shows the data rate per Hz that can be achieved by the secondary user as a function of the utilization of the cellular system. This approach is less effective at 100% cellular utilization than the sharing mechanism described above, although there are still some opportunities to transmit; in this example, the point-to-point link can achieve an average data rate of 0.02 bps/Hz at 100% utilization. No cellular system operates at or near 100% utilization all the time. At 50% utilization, the point to point link can achieve a healthy 0.5 bps/Hz on average with no harmful interference to the cellular system. At any utilization, cooperation leads to slightly better performance.

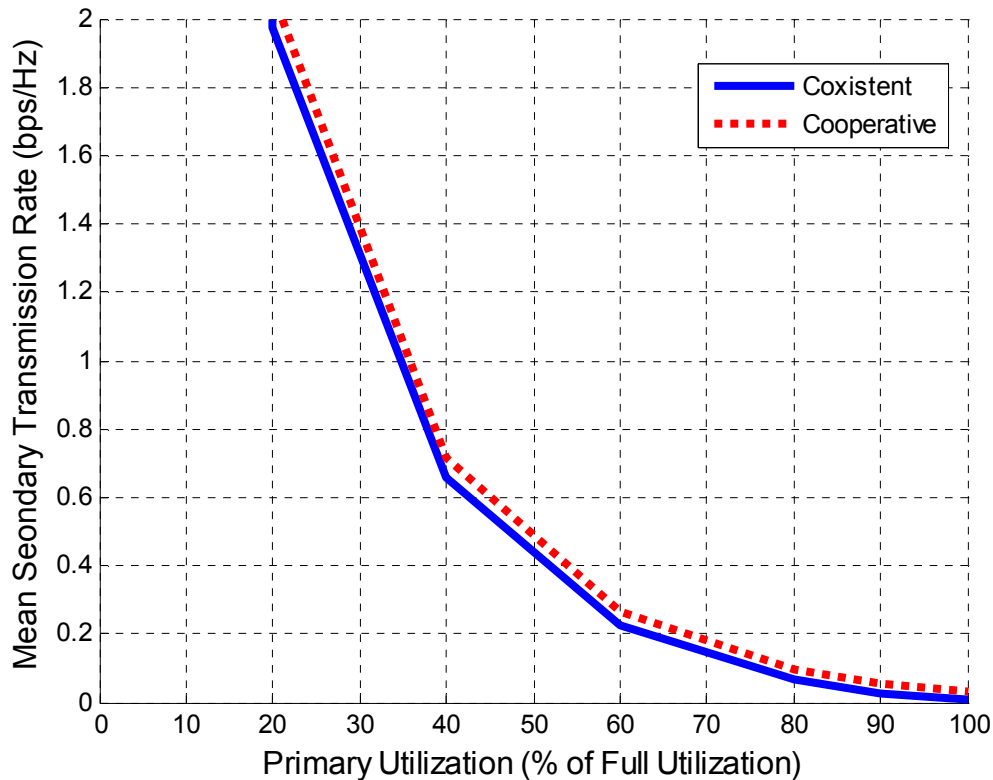


Fig. 1: Mean Transmission Rate of Secondary System vs. Utilization of Cellular System

Other forms of gray-space sharing with cellular are also possible. In (Lee, Han, Hwang & Choi 2009), a secondary device is allowed to transmit in unused time slots of the cellular system, assuming the cellular system uses TDMA technology. Thus, once again, the technology of the secondary system depends on the technology of the primary. In (Bakr, Johnson, Wild & Ramchandran 2008), secondary devices avoid harmful interference to a cellular system by using multi antenna beam forming, combined with knowledge of how the primary system is configured.

Section 2.2: Radar

This section addresses primary-secondary sharing in which radar systems are primary. There is good reason to consider spectrum sharing involving radar. Of the spectrum from 225 MHz to 3.7 GHz, roughly half (1.7 GHz) involves radar or radio-navigation infrastructure (Nebia 2009). There are already proposals (FCC 2011) to make some bands used by radar available for new uses in those safe white spaces where no radar systems are operating, while regions with radar would be “exclusion zones” that are only used for radar. Unfortunately, some of the areas with highest population density, and greatest spectrum scarcity, are likely to be within exclusion zones. There could be great value in opening up

some of these exclusion zones for gray-space sharing, if it could be done without causing harmful interference to radar (Marcus 2009).

This section discusses one example: sharing between fixed rotating radar and a 4G cellular system using technology similar to Long Term Evolution (LTE). (The detailed assumptions and underlying technical analysis can be found in (Saruthirathanaworakun, Peha & Correia 2012).) Radar is the primary user, so the LTE system must be designed such that it never causes harmful interference to radar. Sharing is based on the radar's rotation. From the perspective of an LTE device at a given location, rotation means that the radar's antenna gain fluctuates over time in a pattern that roughly repeats periodically. The LTE device uses dynamic power control to ensure that the interference it causes to the radar falls within the tolerable limit. Depending on the instantaneous value of antenna gain, as well as the LTE device's distance (or more precisely, path loss) to the radar, the device may be prohibited from transmitting, or allowed to transmit at full power, or allowed to transmit at a power below its maximum. This may or may not involve cooperation with the radar. Even in a coexistent model, the calculation of whether an LTE transmission would cause harmful interference to radar requires knowledge of the technical characteristics of the radar, as shown with cellular systems in Section 2.1.

With cell radius held constant, changing LTE power as the radar rotates has the effect of changing LTE data rates. Despite these data rate fluctuations, this shared spectrum can be of great use to a cellular provider. Consider the following scenario. A cellular provider serving a densely populated area where available spectrum is limited uses a combination of dedicated spectrum and spectrum shared with radar. Antennas are collocated, and cells cover roughly the same geographic areas in both the shared and dedicated bands. When capacity is sufficient, all traffic in a given cell is carried in the dedicated band. However, when instantaneous utilization exceeds the capacity of the dedicated band in that cell, some traffic is carried in the shared spectrum. The carrier may choose which traffic to shift to shared spectrum based on quality of service requirements. For example, as shown in (Saruthirathanaworakun, Peha & Correia 2012) and discussed below, voice (telephone) traffic is better left in dedicated spectrum, while video streaming is better shifted to spectrum shared with radar.

Even cells that are fairly close to the radar can support extensive communications on average, although with interruptions and fluctuations in data rate. To quantitatively assess the scheme, we analyzed performance in (Saruthirathanaworakun, Peha & Correia 2012) under the following specific conditions. The radar and cellular systems operate at 2.8 GHz in the same 3 MHz band. (The Federal Aviation Administration operates rotating radar in this band for air traffic control.) It is an urban area with fairly flat terrain, and this is reflected in the signal propagation assumptions (ITU-R 2009; Kurner 1999). The radar rotates at a constant rate, and sends out pulses at a constant power. The radar's rotation period (4.7 seconds), transmit power (0.45 MW (ITU-R 2003)), tolerable Interference-to-Noise Ratio (INR) (-10 dB (Bedford & Sanders 2007)), and peak antenna gain (33.5 dBi (ITU-R 2003)), all of which are static parameters, are known to the secondary system. Background noise is -106 dBm (ITU-R 2003). The radar antenna is a uniformly-distributed aperture type with elevation, azimuthal 3-dB beamwidth, and front-to-back ratio of 4.7°, 1.4°, and 38 dB, respectively (ITU-R 2003; Skolnik 1981). The cellular system uses symmetric Time Division Duplex (TDD), Orthogonal Frequency Division Multiple Access (OFDMA), and

2 × 2 Multiple Input Multiple Output (MIMO) in both directions. At any given time, it employs the modulation (QPSK, 16QAM, or 64QAM) that maximizes data rate under current conditions. To suit an urban environment where demand is high and spectrum is limited, cell radius is 0.8 km. It is assumed that it is sufficiently unusual for a given cell's utilization to exceed capacity of its dedicated spectrum that this rarely happens in adjacent cells simultaneously, so radar is the primary source of interference rather than inter-cell interference. Mobile devices transmit at up to 23 dBm with omni-direction 0 dBi antennas, and basestations transmit at up to 46 dBm with sectorized 18 dBi antennas (Holma, Kinnunen, Kovács, Pajukoski, Pedersen & Reunanen 2009). Background noise spectral density at a secondary receiver is -174 dBm/Hz (Holma, Kinnunen, Kovács, Pajukoski, Pedersen & Reunanen 2009). In this particular scenario, there is only one radar sharing spectrum with an LTE-like cell. Technical analysis will appear elsewhere that explores more complex scenarios with multiple radars and multiple LTE-like cells.

Fig. 2 shows the upstream and downstream data rates achievable in a single LTE cell averaged over the entire rotation period of the radar, as derived in (Saruthirathanaworakun, Peha & Correia 2012). This is a worst-case scenario in which a single mobile LTE device is located at the edge of cell where path loss to the basestation is greatest, and at the point closest to the radar where interference between primary and secondary systems is greatest. Under these numerical assumptions, the radar and cellular basestation would have to be 425 km apart to avoid harmful interference to either system, i.e. to allow both the radar and cellular systems to perform as if they were in dedicated spectrum. Nevertheless, at a distance of just 40 km, the downstream rate of the cellular system is close to what could be achieved in dedicated spectrum on average, although data rate fluctuates. The impact of sharing with radar on the mean achievable upstream rate is somewhat greater, since data rate does not approach the maximum achievable until distance from the radar exceeds 200 km. Thus, this form of sharing can provide somewhat greater spectral efficiency for applications that involve greater downstream data rates. Nevertheless, the extent of communications upstream with relatively small distances to the radar is also impressive. (This is still a simple model, with only one rotating radar operating in the entire band. Forthcoming work will address more complex models, but that is outside the scope of this paper.)

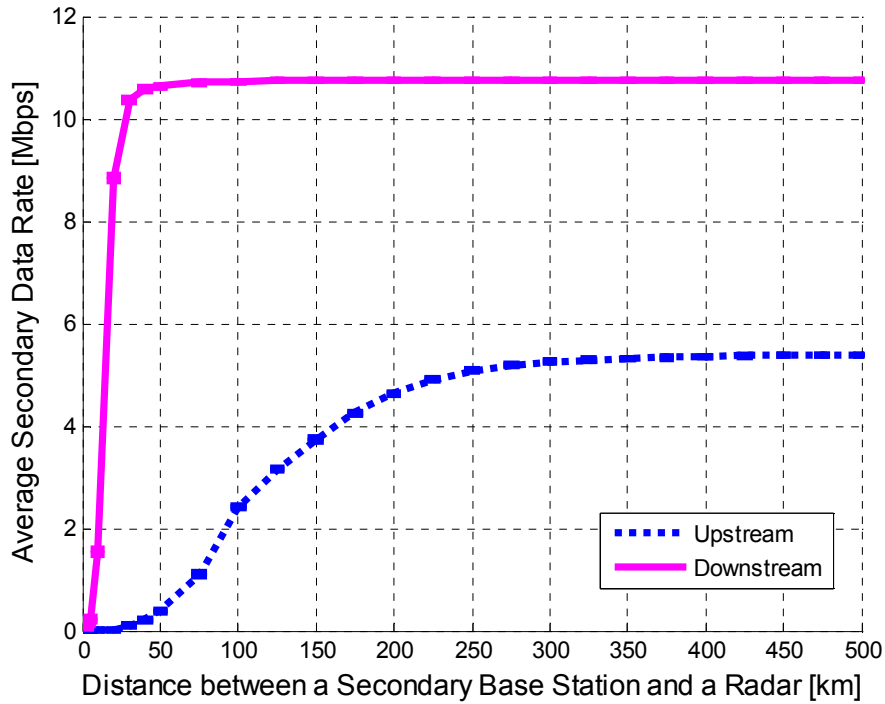


Fig. 2: Mean upstream and downstream data rates of the cellular system

Whether a reasonably high but fluctuating data rate is desirable for a given packet stream depends on that quality-of-service requirements of the associated application. Indeed, whether the fluctuations are noticeable at all depends on the application as well. Consider any application that requires the transfer of a block of data that is useful only when the entire block has been received. If the transfer time equals the rotation time of the radar, or greatly exceeds it, then fluctuations will average out, and data rate as perceived by the application will equal the mean data rate. On the other hand, if transfer time is small compared to radar rotation time, then the time to transfer the file will greatly depend on when in the radar’s cycle the transfer begins. Fig. 3 shows the mean achievable downstream data rate for a file transfer, and the first percentile of perceived data rate² which represents near-worst-case performance, for files of different sizes. (Upstream results are similar (Saruthirathanaworakun & Peha 2010).) When transferring files of multiple MBs, perceived data rate will always be close to the average, but when transferring files of just 1 KB, an application can sometimes perceive data rate to be an order of magnitude below average. Thus, spectrum shared with radar can consistently offer perceived data rates close to those experienced in dedicated spectrum when transferring MBs of information, but not when transferring KBs of information.

² Perceived data rate = (total file length) / [(the time when the file completes transmission) – (the time the file is first available for transmission, regardless of when transmission can actually begin)].

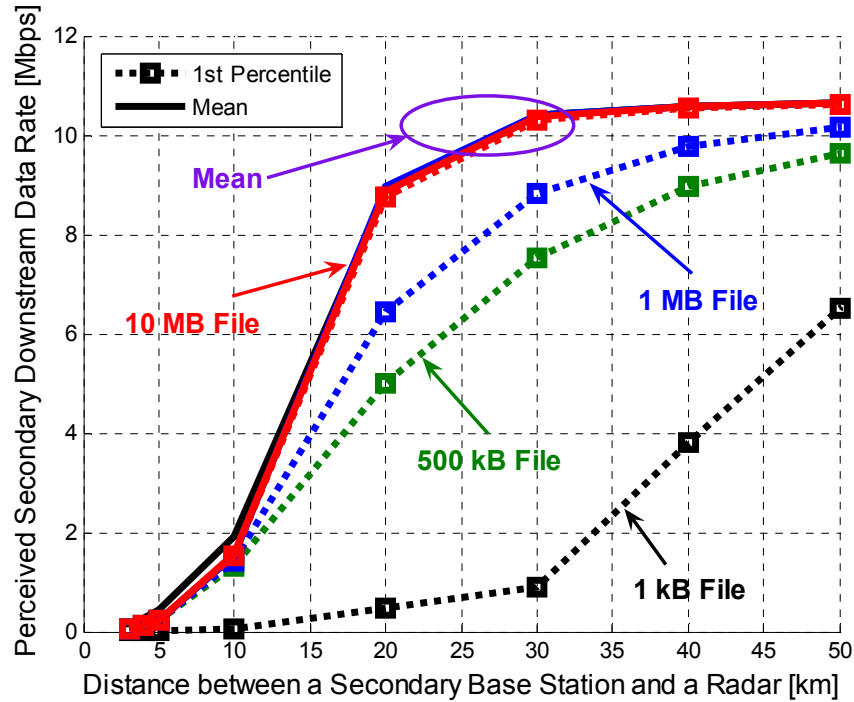


Fig. 3: First Percentile and Mean Perceived Downstream Data Rate vs. Distance between Secondary Base Station and Radar

Based on the above observation, and additional analysis shown in (Saruthirathanaworakun, Peha & Correia 2012), spectrum shared with radar is valuable for a number of common applications. For example, video streaming is a large and rapidly growing portion of Internet traffic (Cisco 2011). It requires a high mean downstream data rate, but with reasonable buffering at the destination (Saruthirathanaworakun, Peha & Correia 2012), the fluctuations in data rate inherent in this approach can easily be tolerated. Web browsing, peer-to-peer file sharing, meter reading, and image transfer can all be well supported in spectrum shared with radar (Saruthirathanaworakun, Peha & Correia 2012). Video, P2P, and web browsing combined are likely to constitute the vast majority of mobile Internet traffic (Cisco 2011). However, any application that requires the transfer of small blocks of information in reliably short periods of time will not be well served in this spectrum. For example, interactive voice over IP (VOIP) is likely to experience unacceptable delays (Saruthirathanaworakun, Peha & Correia 2012); such traffic should be carried in dedicated spectrum.

Section 3: Policies for Gray-Space Sharing

Section 3.1: Gray-Space Technology and Policy Implications

Even where efficient sharing is technically possible and practical, there may be policy challenges to overcome. As shown in Section 2, with gray-space sharing, the gains in spectral efficiency are generally

achieved in ways that make the primary and secondary systems technically interdependent and intertwined. This should lead to technologies, business practices and spectrum policies that differ significantly from what is appropriate with white-space sharing. In a cooperative primary-secondary sharing model, both primary and secondary systems must support the same protocol, so some degree of co-design is a requirement. It is conceivable that a protocol could be developed and standardized that is so general, it applies to a wide range of technologies, thereby allowing primary systems to change without a corresponding change in secondary systems, or vice versa, but this would not be easy to achieve. Although this may be counterintuitive, the interdependence will tend to be even greater with opportunistic primary-secondary sharing models based on coexistence rather than cooperation. In this case, a secondary system must somehow determine from what it can learn about its environment whether a transmission would cause harmful interference, and that typically requires a great deal of information about what interference would be harmful to the specific primary system. This was certainly the case for the cellular and radar examples of coexistent sharing presented in Section 2.

Because gains in spectral efficiency with gray-space primary-secondary sharing often increase the technical interdependence between primary and secondary systems, designers and/or operators of secondary systems must have some means of obtaining detailed information about the technology of primary systems – something that is unnecessary with schemes based on white space sharing. Moreover, with gray-space sharing, if the technology or even the configuration of a primary system changes, then it may be necessary to change the secondary system as well. Some changes might be handled automatically through emerging cognitive techniques, but not all. This problem is less serious for technologies that change very slowly, such as broadcast television, but it could be a big problem for technologies that do, such as cellular. For TV white space devices, designers of secondary systems learn about technical operating requirements from the regulator. To do the same with gray-space sharing, the regulator must either impose technical requirements on the primary system, which would reverse the general trend in spectrum policy towards flexibility (Peha 1998), or must mandate that the primary spectrum user discloses technical information to an extent that would be highly unusual with today's policies. The alternative is to adopt a policy whereby secondary spectrum users must gain access to spectrum through the primary spectrum user instead of the regulator. After all, the primary spectrum user knows its technology, and its plans to change that technology over time. As will be shown in this section, this is often an attractive strategy.

These forms of spectrum sharing also create greater risks of unintended harmful interference than would typically be expected with white space devices, and therefore greater need for new ways to respond rapidly to reports of interference. For example, there may be situations in which a bug causes a secondary transmitter to incorrectly calculate the maximum power at which it can transmit without causing harmful interference to a primary system. Indeed, it is harder to catch bugs like these in secondary systems that adjust to the primary spectrum user, because the bug may only become apparent under circumstances that rarely occur (Peha 2009a). Thus, in addition to an effective process to test sharing mechanisms before deployment, there is greater need for mechanisms with which the secondary system can quickly be forced to reconfigure or discontinue operation if it is ever causing harmful interference in unforeseen ways.

The above issues must be addressed through policy and governance. Some entity must have the ability, the authority, and the responsibility for

- granting or withholding permission to deploy a secondary system in shared spectrum,
- either prescribing or simply relaying information about technical design and operation between primary and secondary spectrum users,
- overseeing an orderly transition when systems are changed or upgraded in a way that could affect any systems with which they share spectrum,
- approving testing procedures intended to ensure that sharing will not cause harmful interference,
- accepting complaints when a system is experiencing harmful interference ,
- identifying the source(s) of harmful interference, and
- requiring that systems be reconfigured or simply turned off sufficiently quickly as needed to end harmful interference

Some of the above capabilities are not needed with white-space sharing, because systems are not so technically intertwined. Others are needed for both white and gray space sharing, but are greatly complicated by the nature of gray-space sharing.

There is a risk that systems run by different organizations will interfere with each other, so both primary and secondary users of spectrum must have confidence in all of above capabilities or they will not invest in systems to be deployed in shared spectrum. In general, with primary-secondary sharing, these functions are undertaken in part by those who hold licenses to be primary spectrum users and in part by a regulator, and the balance of control between these two is a defining feature of the spectrum policy. The best approach depends in large part on whether there is just one primary spectrum user or many, whether there is just one secondary user or many, and how rapidly the relevant technology changes over time. (Note that a spectrum user can be anything from a single transmitter-receiver pair to an entire nationwide infrastructure, provided that it is all under the same administrative control.) The remainder of Section 3 will discuss four very different cases of gray-space sharing, a policy that might effectively support this sharing, and its status under current U.S. spectrum regulation.

Section 3.2: Primary Spectrum Users with Static Technology

Given how gray-space sharing makes systems more technically interdependent, policies can be simpler if technology is static. For sharing that is primary-secondary rather than sharing among equals, the responsibility for avoiding harmful interference is principally borne by the secondary system (Peha 2009a), so what matters is whether the technology of the primary system(s) is static. This is a reasonable assumption for a small number of technologies, such as broadcast radio and television, for which the time between fundamental technology changes is typically measured in decades, and these changes only occur with substantial involvement from the regulator (e.g. the transition from analog to digital TV). This policy has prevailed because a large numbers of devices are controlled by consumers, and these consumers have an expectation that their devices will always work even with competing service providers (e.g. over-the-air TV or radio stations). Thus, the scenario described in Section 1

wherein secondary devices are allowed to transmit in “used” TV spectrum when they are sufficiently close to the TV broadcast tower is an example of gray-space sharing with relatively static technology.

With static technology, some of the capabilities in Section 3.1 that are necessary for gray-space sharing become simple. Regulators already know and publicize the most important technical details of primary systems, and operators of these systems have little cause to complain that proprietary information is being revealed. Regulators are already involved in those rare occasions when the technology of primary systems is transitioned. Nevertheless, spectrum policy must facilitate the other capabilities described in Section 3.1, such as identifying the source of harmful interference to the primary system when it occurs, and shutting down the offending system quickly. If the technology is static, one way to do so involves a variant of the current U.S. policy for white-space sharing in the TV band (FCC 2008) – a policy that puts the regulator rather than primary users of spectrum in charge of controlling access for secondary users. Under current policy, TV white space devices must check a centralized database periodically and cease operation when and where the database requires them to do so. For gray-space sharing, the database might also provide other information, such as the maximum transmit power in a given location, or constraints on how antenna arrays can be configured. Secondary users of this kind could be either unlicensed (as is the case in the TV band today) or licensed. Either way, when there are complaints of harmful interference, the problem can be mitigated quickly by updating the centralized database. Determining which device or devices are causing the harmful interference may be difficult, but because a change in the database effectively stops interference from all devices in a given area regardless of their technology or ownership, it is possible to terminate harmful interference even before the primary source has been identified. Current policy also allows for the possibility that white space devices might someday rely on spectrum sensing instead of a centralized database, although no sensing devices have been approved to date, but gray-space sharing is different. While sensing might someday work well for white space sharing, sensing is more problematic for gray-space sharing because the risk of cumulative unintended interference is higher, and there is no similarly easy way to shut off the myriad unlicensed devices when problems occur, or even when technology changes.

Although allowing gray-spacing sharing of TV bands is conceptually similar to current U.S. policies for white-space sharing, current policy prohibits gray space sharing in the TV band (and obviously should do so unless and until a specific sharing technology has been shown to prevent harmful interference with sufficient reliability).

Section 3.3: A Single Primary Spectrum User

Where technology is dynamic, the situation is simplest when there is a single primary license-holder for all of the spectrum that would be shared with secondary devices. In this case, the license-holder is well-positioned technically to provide all of the capabilities listed in Section 3.1. Thus, there is good reason to consider a policy whereby the license-holder rather than the regulator has the authority to determine whether, when, and how another system may access the spectrum on a secondary basis. Unlike the regulator, the license-holder has detailed knowledge of its current technology and all plans for upgrades, which it can share with secondary users when (and perhaps only when) this facilitates gray-

space sharing. The license-holder also has the ability to detect harmful interference to the only primary user (itself) much faster than a regulator, and the ability and motivation to promptly demand that secondary systems cease causing harmful interference if it does occur. A requirement for secondary users to react immediately can be written into the contract between each secondary user and the license-holder. Unlike a regulator, a license-holder's goal is not to improve overall utilization of spectrum resources, but a license-holder could have incentive to make its spectrum available for gray-space sharing if it is allowed under spectrum policy to charge secondary spectrum users. Payments may take the form of usage-based micropayments (Peha & Panichpapiboon 2004) for the right to transmit briefly in "used" spectrum, or long-term subscriptions paid by the users of the device, or subscriptions for the lifetime of the device paid by the manufacturer or importer of the device (Peha 2009a).

In the U.S., this form of sharing is currently legal in many but not all bands under the FCC's secondary market rules (FCC 2003, 2004) which give license-holders in many bands the right to grant secondary users access to their spectrum. There is currently a requirement that the license-holder notify the FCC about these transactions afterwards, but otherwise, the regulator plays no direct role in this kind of sharing. However, gray-spacing sharing does not fit the secondary spectrum market that many people imagine, wherein "blocks" of spectrum with clear frequency and geographic boundaries are exchanged for extended periods. In gray-space sharing, different systems operate in the same block of spectrum, and users must agree to more technical details than what happens at a boundary between adjacent blocks. New legal and business practices probably must emerge and become reasonably standardized before such arrangement can be made at low transaction costs.

Secondary market rules do require secondary spectrum users to operate in compliance with any restrictions on the license-holder. This is not a problem in bands which offer complete flexibility (Peha 1998; Peha 2009a) of technology and use, such as cellular bands. However, in cases where the license restricts technology and transmitter location, secondary market rules would not allow the license-holder to make gray-space sharing arrangements. For example, if a license only allows satellite communications, secondary market rules would not allow gray-space sharing between satellite and terrestrial communications. Similarly, gray space sharing with TV of the kind discussed in Section 1 would not be possible under secondary market regulations because TV licenses give broadcasters little flexibility.

Section 3.4: A Single Secondary Spectrum User

The next simplest case is the one with a single secondary user, but devices operated by that secondary user might simultaneously interfere with multiple primary spectrum users. For example, this might occur in the scenario presented in Section 2.2 where a single cellular provider controls the cells throughout a region, but there are multiple radar systems with primary rights, and each radar is operated by a different government agency or commercial company. It also might occur when secondary users are broadcasters.

The fact that there is a single secondary spectrum user simplifies some of the capabilities described in Section 3.1, but not all. In particular, when there are complaints of harmful interference, it is clear who to hold accountable, and easier to require that portions of the presumably offending system be modified or shut down quickly. However, how will the designers or users of secondary systems obtain sufficiently detailed information about the current technology used by primary systems to avoid causing harmful interference? If technology is static, a regulator can easily provide or oversee this spectrum sharing when granting secondary users access to spectrum, perhaps with a policy similar to the one in Section 3.2, but this is problematic when technology is heterogeneous across primary spectrum users, changing over time, or both.

This kind of gray-space sharing might be supported through a two-step process. First, a single secondary spectrum user gains exclusive rights by obtaining a license for secondary access over a region that includes multiple primary spectrum users. For example, a spectrum band with many government systems may be auctioned for commercial use with a requirement not to cause harmful interference to incumbents. Initially, the area around these incumbent systems plus a buffer zone could be “exclusion zones” that are entirely off limits for secondary use, making this a form of white space sharing. Then, under secondary market rules (FCC 2003, 2004), the secondary user must bargain with each primary spectrum user to increase spectral efficiency through gray-space sharing, and this bargaining will include sharing detailed technical information. Such bargaining may include adjustments in both the transmitters and receivers used by both primary and secondary systems (Franzoni and Crocioni 2011). However, these are not simply negotiations over what happens at the frequency or geographic boundary between the spectrum blocks used by each system, because with gray-space sharing, they occupy the same spectrum and there are no clear boundaries.

Section 3.5: Multiple Primary Users, Multiple Secondary Users

In the most general case, new policies and governance structures are needed to facilitate gray-space sharing. Consider the radar example in Section 2.2, except with many radar systems, each operated by a different organization, and possibly using a different technology. Each radar operator must be concerned about the cumulative interference from multiple secondary systems, and each secondary system must be concerned about the interference it might cause to multiple radars. Thus, there is no single party that can determine where it is safe to deploy a new secondary system. Moreover, an upgrade to a single radar may change where and how a number of secondary transmitters can operate without causing harmful interference. Regulators are not set up well to manage the complexity of preventing interference among diverse and changing systems that are close enough together to potentially cause interference. Worse yet, under current U.S. policy, if this sharing would involve both government radar operators and commercial entities, it would probably involve two different federal agencies: the Federal Communications Commission and the Department of Commerce. This is even more problematic. Thus, there are serious impediments to gray-space sharing among systems under different administrative control.

New governing bodies are needed to provide the capabilities described in Section 3.1. In effect, such a governing body would serve as a band manager for a specific block of spectrum, but would be limited to

allowing uses of the band that never cause harmful interference to a primary user. This governing body must inspire confidence in primary spectrum users, which are operators of radar in this case, that harmful interference will not be a problem. This governing body might include representatives from organizations that operate systems in the band on a primary basis. If the number of primary systems in the band is large, responsibility may be divided regionally. At the same time, interference protection cannot be the sole goal of the governing body, or it will never be motivated to allow sharing. There must also be incentive to expand use of the spectrum. Finding incentives for primary spectrum users in situations with multiple primary and secondary systems is an open and challenging problem. Finally, this governing body must inspire confidence in secondary spectrum users that, even though their rights are secondary, those rights will be sufficiently stable over time to warrant long-term investment. Thus, the body must be given substantial authority to make decisions that are binding on primary license-holders rather than merely advisory. An approximate model for this body are the Regional Planning Committees (RPCs) in the U.S., which coordinate spectrum use in a given geographic region in the bands allocated to public safety agencies so as to minimize mutual interference. Nevertheless, no body yet exists that can address gray-space sharing more generally.

Section 4: Conclusions

While much of the discussion of increased spectrum sharing has focused on deploying devices in frequency bands and regions in which no wireless devices are deployed, there are also opportunities to share spectrum that is being used, which we refer to here as “gray-space” sharing. Indeed, it is often possible to deploy additional devices on a secondary basis in spectrum that current users would view as 100% utilized, without causing any harmful interference to the primary spectrum users. Through examples, this paper has shown that a great deal of communications among secondary devices is technically possible in spectrum that is also heavily used by cellular carriers or radar or television.

A sharing arrangement that is technically possible still requires the right forms of policy and governance. Compared with white space sharing, gray space sharing forces the systems that share spectrum to become more technically interdependent. This can require more information sharing among different types of users and equipment designers, and it can make the process of upgrades far more complex. This makes spectrum-sharing policies that do not require much interaction between primary and secondary spectrum users less appropriate for gray space sharing, such as the TV white space approach (FCC 2008), or the policy described in Section 1 where secondary users can essentially be squatters in blocks of spectrum that are clearly idle.

In cases where a single entity is licensed to be the primary spectrum user over a large area (i.e. large compared to the size of the system deployed by a secondary user), secondary market rules (FCC 2003, 2004) provide a useful basis for sharing in many cases, although these markets would lease out opportunities to transmit at a much smaller granularity than has emerged to date. Moreover, the technical parameters underlying gray-space sharing are qualitatively different from what many envision for spectrum markets, which may require the creation of new legal and business arrangements. In this

model, the license-holder is in control. Where there are many primary systems, an alternative approach is to grant exclusive secondary rights to one spectrum user, and allow that user to bargain for gray-space sharing with each primary system on a bilateral basis. Again, each primary license-holder is in control. In contrast, in cases where technology is relatively static, as occurs in broadcasting, there are scenarios in which the regulator rather than a license-holder may be able to make spectrum available for secondary use. One method to do so involves an extension of the approach developed for TV white space devices.

However, in cases where there are multiple systems with primary rights under different administrative control, and technology is not static, new policies are needed to facilitate this form of sharing. One example is where multiple radar systems operate in a band, and would share that band with other systems on a primary-secondary basis. For a case like this, a new governing body should be established to facilitate sharing in the band, on a regional or national basis. This body must have the ability, authority, and responsibility for introducing sharing while protecting primary users. This includes ensuring that technical information is shared where necessary, determining where sharing is possible, managing technology changes over time, and addressing any incidents of harmful interference that do occur.

Spectrum managers around the world are seeking ways to make spectrum available, and this may create new opportunities for some of the gray-space sharing approaches described in this paper. For example, the U.S. Department of Commerce will soon make a recommendation regarding spectrum at 1750 MHz, a band that is currently used by many government agencies. Freeing entire swaths of this spectrum will require moving many incumbents, and some moves may be expensive and time-consuming. Allowing white-space and gray-space sharing may be an attractive alternative in some bands.

Acknowledgements

My thanks to Rathapon Saruthirathanaworakun, a Carnegie Mellon University Ph.D. student who works with the author, and who did outstanding work on some of the technical analysis described in this paper (Saruthirathanaworakun & Peha 2010; Saruthirathanaworakun, Peha & Correia 2012a; Saruthirathanaworakun, Peha & Correia 2012b).

References

- [1] Bakr, O., Johnson, M., Wild, B. & Ramchandran, K. (2008). A multi-antenna framework for spectrum reuse based on primary-secondary cooperation. *Proceedings of IEEE DySPAN*.
- [2] Bedford, B. & Sanders, F. (2007). Spectrum sharing and potential interference to radars, *Proceedings of the International Symposium on Advanced Radio Technologies (ISART)*.
- [3] Cisco (2011). *Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2010–2015*, White Paper, Cisco, USA, Feb. 2011.

www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf

- [4] De Vany, A., Eckert, R., Meyers, C., O'Hara, D. & Scott, R. (1969). A Property System for Market Allocation of the Electromagnetic Spectrum: A Legal-Economic-Engineering Study. *Stanford Law Review*, 21(6), 1499-1561.
- [5] Federal Communications Commission (2002), Spectrum Efficiency Working Group, Technical Report, WT Docket No. 02-135, Spectrum Policy Task Force Report, Nov. 2002.
- [6] Federal Communications Commission (2003). Promoting Efficient Use of Spectrum Through Elimination of Barriers to the Development of Secondary Markets, Report And Order And Further Notice Of Proposed Rulemaking, WT Docket No. 00-230, Oct 2003.
http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-03-113A1.pdf
- [7] Federal Communications Commission (2004). Promoting Efficient Use of Spectrum Through Elimination of Barriers to the Development of Secondary Markets, Second Report And Order, WT Docket No. 00-230, Sept.2004. http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-04-167A1.PDF
- [8] Federal Communications Commission (2008). In the matter of unlicensed operation in the TV broadcast bands, 2nd Report and Order and Memorandum Opinion and Order. ET Docket No. 04-186. Nov 14, 2008. www.wispa.org/wp-content/uploads/2008/11/FCC-08-260A1%282ndRO-111408%29.pdf
- [9] Federal Communications Commission (2010). *The National Broadband Plan: Connecting America*. www.ece.cmu.edu/~peha/policy.html
- [10] Federal Communications Commission (2011), Spectrum Task Force Requests Information on Frequency Bands Identified by NTIA as Potential Broadband Spectrum, Public Notice, FCC DA 11-444, March 2011. http://transition.fcc.gov/Daily_Releases/Daily_Business/2011/db0308/DA-11-444A1.pdf
- [11] Franzoni, P & Crocioni, L. A. (2011) Transmitters and Receivers' Investment to Avoid Interference: Is There an Optimal Regime?, *Telecommunications Policy*, 35(6), 568-578.
- [12] Holma, H., Kinnunen, P., Kovács, I.Z., Pajukoski, K., Pedersen, K. & Reunanen, J. (2009). Performance," in Holma H. & Toskala, A. (Eds.), *LTE for UMTS: OFDMA and SC-FDMA Based Radio Access*, John Wiley & Sons, UK.
- [13] ITU-R (2003). *Recommendation M.1464-1: Characteristics of radiolocation radars, and characteristics and protection criteria for sharing studies for aeronautical radionavigation and meteorological radars in the radiodetermination service operating in the frequency band 2 700-2 900 MHz*.
- [14] ITU-R (2009). *Recommendation P.1546-4: Method for point-to-area predictions for terrestrial services in the frequency range 30 MHz to 3 000 MHz*.
- [15] Kurner, T. (1999). Propagation Model for Macro-Cells, in Damosso, E. & Correia, L.M. (Eds.), *Digital Mobile Radio towards Future Generation Systems: COST 231 Final Report*, European Communities, Belgium.
- [16] Lee, H., Han, K., Hwang, Y. & Choi, S. (2009). Opportunistic band sharing for point-to-point link connection of cognitive radios. *Proceedings of IEEE International Conference on Cognitive Radio Oriented Wireless Networks (CrownCom)*.

- [17] Marcus, M. J. (2009). Sharing government spectrum with private users: opportunities and challenges, *IEEE Wireless Communications Magazine*, 16(3), 4-5.
- [18] National Research Council (2010). *Spectrum Management for Science in the 21st Century*, National Academies Press.
- [19] Nebia, K. (2009). US National Telecommunications and Information Administration (NTIA), *Presentation: spectrum with significant federal commitments, 225 MHz - 3.7 GHz*.
www.ntia.doc.gov/advisory/spectrum/csmac_2009_december_meeting_reports.html
- [20] New America Foundation and Shared Spectrum Company (2003). Dupont Circle Spectrum Utilization During Peak Hours, Technical Report.
- [21] Panichpapiboon, S. & Peha, J. M. (2008). Providing Secondary Access To Licensed Spectrum Through Coordination, *ACM/Baltzer Wireless Networks*, 14(3) 295-307.
www.ece.cmu.edu/~peha/wireless.html
- [22] Peha, J. M. (1998). Spectrum Management Policy Options, *IEEE Communications Surveys*, Fourth Quarter 1998. www.ece.cmu.edu/~peha/wireless.html
- [23] Peha, J. M. (2009a). Sharing spectrum through spectrum policy reform and cognitive radio, *Proceedings of the IEEE*, 97(4), 708-19. www.ece.cmu.edu/~peha/wireless.html
- [24] Peha, J. M. (2009b). A Spectrum Policy Agenda, In A. M. Schejter (Ed.). ... *and Communications for All: A Policy Agenda for the New Administration*, Lexington Books.
- [25] Peha, J. M. & Panichpapiboon, S. (2004). Real-Time Secondary Markets for Spectrum, *Telecommunications Policy*, 28(7-8), 603-18. www.ece.cmu.edu/~peha/wireless.html
- [26] Saruthirathanaworakun, R. & Peha, J. M. (2010). Dynamic primary-secondary spectrum sharing with cellular systems," *Proceedings of IEEE International Conference on Cognitive Radio Oriented Wireless Networks (CrownCom)*. www.ece.cmu.edu/~peha/wireless.html
- [27] Saruthirathanaworakun, R. , Peha, J. M. & Correia, L. M. (2012). Opportunistic Primary-Secondary Spectrum Sharing with a Rotating Radar, to appear in *Proceedings of IEEE International Conference on Computing, Networking and Communication*. Extended version to appear as Opportunistic Sharing between Rotating Radar and Cellular in *IEEE Journal on Selected Areas in Communications*. www.ece.cmu.edu/~peha/wireless.html
- [28] Shared Spectrum Company (2004). New York City Spectrum Occupancy Measurements September 2004.
- [29] Skolnik, M.I. (1981). *Introduction to Radar Systems*, International Ed., McGraw-Hill, Singapore.