

5-2004

Real-Time Secondary Markets for Spectrum

Jon M. Peha

Carnegie Mellon University, peha@andrew.cmu.edu

Sooskan Panichpapiboon

Carnegie Mellon University

Follow this and additional works at: <http://repository.cmu.edu/epp>



Part of the [Engineering Commons](#)

Published In

Telecommunications Policy, 28, 603-618.

This Article is brought to you for free and open access by the Carnegie Institute of Technology at Research Showcase @ CMU. It has been accepted for inclusion in Department of Engineering and Public Policy by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Real-Time Secondary Markets for Spectrum

Jon M. Peha* and Sooksan Panichpapiboon†

Carnegie Mellon University

Abstract

Spectrum licensing is one effective way to guarantee adequate quality of service for license-holders, because they can have exclusive access to spectrum. However, exclusivity also leads to inefficient use of spectrum, which contributes to today's shortage of spectrum. Unlicensed spectrum is a valuable way to increase sharing, and potentially achieve much greater spectral efficiency, but guaranteed quality of service is not available in unlicensed spectrum, which is a serious problem for some applications. This paper explores a regulatory alternative where license-holder and secondary spectrum users share spectrum, thereby increasing spectral efficiency and alleviating spectrum scarcity, and where quality of service can be guaranteed. This is a *real-time secondary market*, where secondary users ask the license-holder for temporary access to spectrum as needed. The license-holder permits this sharing when and only when it determines that quality of service requirements can still be met for both license-holder and secondary users. In this paper, we quantitatively assess the costs and benefits of a real-time secondary market for the special case where the license-holder happens to be a GSM-based cellular carrier. We demonstrate that the license-holder can let many secondary users share spectrum with little impact on the capacity available to cellular customers. We further show that the cellular carrier profits from this arrangement even if the price for secondary access is quite low, making this an attractive scheme for both license-holder and secondary users. Finally, we address the challenges of transferring funds from secondary user to license-holder, and show that solutions to these challenges do exist.

1 Introduction

There has been much discussion of an alleged shortage of spectrum, but simple scanning shows that much of the usable spectrum is essentially idle at any given time. This inefficiency

*Professor of Electrical Engineering and Public Policy at Carnegie Mellon University (CMU), Associate Director of CMU Center for Wireless and Broadband Networks, and Chief Technical Officer of Cyphermint, peha@cmu.edu, www.ece.cmu.edu/~peha and www.cyphermint.com

†Graduate Student at Carnegie Mellon University Department of Electrical and Computer Engineering, sooksan@andrew.cmu.edu

is the legacy of spectrum regulations designed for old technology. Regulatory reform could unleash new technology that would help to alleviate spectrum scarcity. This paper focuses on the potential benefits of allowing real-time secondary markets for spectrum, which may be the next logical step beyond the secondary markets recently considered by the US Federal Communications Commission (FCC) [1–3]. These markets would fulfill a need not met by other common regulatory paradigms, such as traditional licensing and unlicensed operation.

Whenever someone who needs a resource sporadically is given exclusive access to that resource, some inefficiency is inevitable. This occurs with traditional licensing. For example, if a cellular carrier has a license which guarantees exclusive access to a block of spectrum, then some capacity will sit idle when the number of calls underway requires less than full capacity, which typically occurs more than 98% of the time. It has been shown quantitatively that dividing a given block of spectrum into separate exclusive blocks greatly reduces the amount of traffic (e.g. cellular telephone calls per minute) that can be carried in this block [4, 5]. Of course, exclusive access does offer an advantage that is essential for some applications: guaranteed quality of service. Indeed, licensing began so that a radio broadcaster or other spectrum-user could build infrastructure without fear that others would later interfere with its transmissions.

Unlicensed spectrum offers the opposite traits. It forces spectrum sharing, thereby opening up the possibility of much greater spectral efficiency. This efficiency is not guaranteed, since users sharing a resource have less incentive to conserve it, and there are circumstances where this can lead to tremendous inefficiencies. However, there are technical and regulatory ways to address this problem, including an effective etiquette [6–10]. Giving users the ability to deploy and move transmitters anywhere at any time without first getting a regulator’s permission has many other important advantages as well. Nevertheless, there is no limit to the number of devices that may be simultaneously competing for unlicensed spectrum, so there is no way to guarantee quality of service, as is possible in licensed spectrum.

Different applications deserve different regulatory schemes [11]. Broadcast television is well served with traditional licensing; the broadcasters want guaranteed quality of service, and should be more than willing to pay the cost of licensing to get it. Moreover, transmissions are continual at a constant data rate so exclusivity does not yield such inefficiency. On the other hand, consumers who want (sporadic) broadband access to the Internet may be willing to surrender their quality of service guarantees, so they can connect to their local ISP at little cost over unlicensed spectrum. However, neither the typical licensed or unlicensed approaches quite fit an application that uses spectrum sporadically and requires guaranteed quality of service. For example, consider a television station that occasionally carries a live news event and needs a wireless connection from the portable camera at the event back to the studio. This station cannot risk interference during a live broadcast, so guaranteed quality of service is important. However, giving the station exclusive access to spectrum for this purpose is highly inefficient and costly, because the spectrum would sit idle much of the time. As another example, there are Internet users who prefer and are willing to pay for guaranteed quality of service on the connection to their ISP (as DSL would provide), so they may not be happy with access over unlicensed spectrum. For applications like these,

we propose real-time secondary markets, which offer both the efficiencies of sharing and the possibility of quality of service guarantees.

In this scheme, spectrum is licensed. Secondary devices gain the right to transmit by explicitly requesting permission from the license-holder as needed. The license-holder grants permission only if doing so will not cause excessive interference with others calls that are already underway, and if the license-holder can promise the secondary device that the latter's quality of service requirements will be met, despite any transmission from the primary license-holder and from other secondary devices. The license holder can charge a fee for this, which gives license-holders incentive to share the spectrum, and it gives secondary users incentive not to waste that spectrum.

Secondary access, where a secondary device is allowed to transmit if and only if it does not interfere with the primary license-holder, is far from new. A secondary market is different. With traditional secondary access, secondary users get permission to transmit from the regulator, and not the license-holder. Moreover, the secondary user usually gets no quality of service guarantees, because a primary user may need the spectrum. With a secondary market, guarantees are possible, through explicit coordination between license-holder and secondary users.

The US Federal Communications Commission has lately been considering the possibility of secondary markets where license-holders negotiate with secondary users [1, 2]. Very recently [3], they announced that some forms of secondary markets will be allowed. This is an important step towards the vision described in this paper. However, the secondary market addressed in those FCC proceedings does not approach real time: a license-holder is expected to notify the FCC 21 days before a secondary user begins operation. Such rules may be appropriate if a license-holder wishes to lease a block of spectrum to a PCS carrier for five years, but not if it wishes to lease a block to a television news show for five minutes. Indeed, the FCC defined a lease to be "short term" if it is under 360 days, so they apparently were not thinking of leases lasting minutes or seconds. The work described in this paper may help the FCC and other regulatory and legislative bodies decide whether to extend the secondary market approach to short-term real-time deals.

Any license-holder could consider making spectrum available in a real-time secondary market, regardless of what the license-holder uses that spectrum for. In the special case where the primary license-holder does not use the spectrum at all, and only makes it available to secondary devices, the secondary market considered here becomes identical to a *band manager* scheme [11, 12]. One challenge with the band manager approach when real-time access is required is that a secondary device must be able to communicate with the band manager whenever spectrum is desired, perhaps over a wireless link. This may force the band manager to build costly infrastructure so there is always a transceiver near any secondary device. Some license-holders already have that infrastructure in place, such as a cellular carrier. In this paper, we consider the case where the primary license-holder is a GSM-based cellular carrier that wants to provide access to spectrum in a real-time secondary market.

One objective of this work is to quantitatively assess the potential impact on spectrum

utilization of this kind of spectrum sharing. We show that letting other devices share the spectrum has only a small impact on the performance of the license holder, and that this sharing scheme is viable technically and economically.

Two technologies have emerged in recent years that make real-time secondary markets more attractive, although neither is actually required. The first is inexpensive global positioning system technology (GPS), with which devices can determine their approximate location. A license-holder can use this location information to estimate when two devices might interfere with each other before allowing a secondary device to transmit, which can be used to greatly increase frequency reuse. The second emerging technology is software-defined radio, which will make it possible for relatively inexpensive devices to move from one frequency band to another. Thus, if a secondary device cannot currently get the desired quality of service guaranteed in one frequency band, the device can jump to another band [13].

New Internet payment systems also constitute an enabling technology, by providing an easy way for secondary users to pay the license-holder. The payment mechanism should allow a secondary user to buy services while communicating only with the license-holder, and without creating opportunities for fraud.

Section 2 provides a general description of our scheme. We then discuss the precise conditions for call admission in Section 3. Section 4 quantitatively describes the impact of secondary users on cellular customers, and related economic implications. Section 5 addresses the payment mechanism. Finally, we conclude this paper in Section 6.

2 Basic Model

A primary user is any license holder that allows its spectrum to be shared by other devices at its own discretion. Here, we study the case where the primary user is a GSM-based cellular network operator. We assume that the cellular carrier is capable of locating the position of each primary device, which is consistent with the FCC's E911 requirement [14,15]. A number of techniques for locating a mobile device in a GSM network have been proposed [16,17]. In this paper, we also restrict ourselves to the simpler case where the base station in one cell does not coordinate with other cells. Coordination could improve performance, so this is a conservative assumption, to be relaxed in future work.

In this paper, we make two assumptions for secondary devices. First, secondary users only use downstream channels of the cellular network. This is also a conservative assumption. Second, each secondary device has GPS capability. Some examples of applications that may be appropriate for secondary access are broadband wireless last-mile or middle-mile internet access, a microwave link, and any point-to-point communications that requires a guarantee of quality of service.

The cellular base station must communicate with all secondary devices in the cell that would like to transmit. The base station would obtain the following information: the transmit power, the amount of bandwidth needed, the amount of tolerable interference, and the

location. When communicating with a cellular network, a secondary user can use the signaling channels already used by the cellular network or it can use an unlicensed band.

The secondary transmitters know all of the above information, with the possible exception of the location of the secondary receivers. There are several ways for a base station to obtain this information. The secondary transmitter could request the location of the secondary receiver by communicating over unlicensed spectrum and then the secondary transmitter could include both locations when making its request to the cellular base station. Alternatively, the secondary transmitter and the secondary receiver could independently register with a cellular base station. To handle the case where the secondary transmitter and the secondary receiver are in different cells, each primary base station would have to share this location information with other base stations.

3 Admission of Primary Calls

Depending on their locations, primary and secondary device can sometimes coexist in the same channel of the same cell. When this is possible, the secondary device has no impact on cellular users whatsoever. We begin analysis by defining precise conditions necessary for a primary call to be admitted, and thus by determining when this form of frequency reuse is possible. (See [18] for analogous discussion of when a secondary call can be admitted.)

In admitting a new cellular call, the base station must make sure that the signal to interference ratio (SIR) of the new primary call and the SIR of the secondary calls underway are adequate. We will assume that a base station uses a power control scheme, so the power transmitted by the base station will vary with the location of the primary device, and the power of the received signal will be constant, independent of location.

A blocking area is defined as an area where primary devices are not able to share a channel or channels with ongoing secondary calls. This can happen when the primary device is too close to existing secondary users, and it can also happen when the distance between the primary device and the base station is so great that the base station would have to transmit at a higher power than ongoing secondary calls could tolerate. The following analysis derives the minimum separation required between a primary handset and a secondary user, and the maximum distance that a primary handset can be from the base station. We will later use these two distances to calculate the size of a blocking area.

Condition 1: SIR on the primary downlink must be adequate.

Let d_{min} be a minimum distance required between a primary handset and a secondary transmitter. Assuming that signal attenuates with distance to the power of n and perfect power control, the signal strength on the primary downlink is P_{max}/C^n , independent of the location of primary device, where C is a cell radius, P_{max} is the maximum transmit power

of a base station, and n is a path loss exponent. Thus, d_{min} can be written as

$$d_{min} = C \left(\frac{X \cdot P_{ts}}{P_{max}} \right)^{1/n} \quad (1)$$

where X is the SIR required on primary downlink and P_{ts} is the transmit power of a secondary device.

Condition 2: SIR of an ongoing secondary call must be adequate.

Let d_{max} be the maximum distance a primary handset can be from the base station. Using the same assumptions as in the previous case, d_{max} can be written as

$$d_{max} = d_2 \cdot C \left(\frac{Z}{P_{max}} \right)^{1/n} \quad (2)$$

where d_2 is a distance between a secondary receiver and the base station, Z is the tolerable interference of a secondary user, and the other variables are as previously defined.

The blocking area created by a secondary user is the union of a circle with radius d_{min} centered at a secondary transmitter and the area outside the circle with radius d_{max} centered at the base station.

In practice, there is often uncertainty in the location of devices, so the measured distance between two devices may not equal the actual distance. To be conservative, a device is in the blocking area if

- measured distance between a device and the closest secondary transmitter $< d_{min} + E_{95}$,
or
- measured distance between a device and the closest base station $> d_{max} + E_{05}$,

where E_{05} and E_{95} are the 5th and 95th percentile of the error in the distance measurement, respectively (as calculated in [18]). (Note that generally $E_{05} \leq 0$ and $E_{95} \geq 0$.)

4 Impact of Secondary Users on Quality of Service

In this section, we quantify the impact of allowing secondary devices to use the spectrum licensed to the cellular network. We evaluate the extent to which carriers would be forced to serve fewer cellular customers, and whether carriers would have financial incentive to try this form of spectrum sharing.

We assume that secondary devices are stationary during periods when they are using spectrum. In addition, it is assumed for this analysis that call holding time of a secondary call is much longer than that of the primary call. Consequently, the transient period after

the secondary call begins or ends can be ignored when calculating blocking probability of the cellular calls, so steady-state analysis can be used. We calculate blocking probability of cellular calls as if primary devices are fixed during the calls. We assume that call arrivals of the primaries follow a Poisson process, call holding time is exponentially distributed, and primary calls are uniformly distributed throughout the cell. We also assume that an *internal handoff*, which is a handoff within a cell, is allowed. In other words, a primary call can be moved from one channel to another at any time. Secondary devices are in randomly selected locations, and we assume here that the location distributions of these transmitters are uniform and independent. Results are achieved through simulation. (See [18] for closed-form analysis of the case where there is only one secondary call per cell at a time.)

For the simulation results in this section, we use the following values unless stated otherwise, all of which are reasonable for a GSM network. The maximum transmit power of a base station is 30 W. Cell radius is 2 km. Path loss exponent is 3.5. The number of 200-kHz channels in the cell is 41. The transmit power of each secondary device is 100 mW, and each secondary call requires 200 kHz of bandwidth. The SIR required by a cellular call and a secondary call is 15 dB.

Section 4.1 introduces key metrics for measuring system performance. In Section 4.2, we show system performance, as derived with two different frequency assignment algorithms that will be used in subsequent subsections. Section 4.3 evaluates the impact of inaccuracy in location measurements. Finally, in Section 4.4, we quantitatively evaluate the economic implications of secondary users.

4.1 Metrics of Performances

We will use three kinds of metrics in measuring how our sharing scheme affects the cellular network: the expected call blocking probability of cellular calls, admissible traffic load of cellular calls, and the *minimum break even rate*, which is the per-minute rate that a carrier needs to charge each secondary call in order to generate the same revenue as in the case where there is no sharing.

We calculate the expected call blocking probability by repeatedly running the simulation with the same number of secondary calls in different locations and averaging the blocking probabilities collected in all runs. In each run, a series of locations for secondary calls is randomly selected. When a location is selected for which all channels are unavailable, a new location is randomly selected. This is repeated until the desired number of secondary calls in the cell is reached. In all the graphs presented, the expected call blocking probability has a 95% confidence interval that is within $\pm 5\%$ of the value shown.

Admissible traffic load is the amount of traffic that a network can carry while maintaining a constant expected call blocking probability. In all the graphs presenting admissible traffic load, the admissible traffic load shown yields an expected call blocking probability between 1.95% and 2.05%. Each expected call blocking probability is also accurate within $\pm 5\%$ with a 95% confidence interval.

For a cellular carrier to gain the same revenue as in the case where there is no spectrum sharing, the following equation must hold.

$$\frac{S}{P} = \frac{(1 - B)(A_0 - A)}{N} \quad (3)$$

where S is the minimum break even rate, P is a per-minute rate that a carrier normally charges a cellular call, B is a blocking probability of a cellular call, A is admissible traffic load of cellular calls, A_0 is admissible traffic load of cellular calls when no secondary calls are present, and N is the number of secondary calls in the cell. We will refer to the ratio in (3) as the “normalized minimum break even rate.”

4.2 Performance of Frequency Assignment Algorithms

The effect of secondary devices on cellular call blocking probability depends in part on the *frequency assignment algorithm*, which determines the frequency at which a secondary device will operate. We introduce two simple schemes which are “First Available” (FA) and “Load Balancing” (LB). In the FA scheme, the network operator always scans the channels in the same order, and the first frequency found where interference is within the acceptable limit is assigned. In the LB scheme, in contrast, the network operator tries to spread the number of secondary users equally on different frequencies by scanning channels in the order determined by the number of secondary calls underway, beginning with the frequency where the number of secondary calls is smallest.

It can be observed from Fig. 1 that when there are few secondary calls in the cell, the LB scheme performs better than the FA scheme. This is because when there are not many secondary calls, spreading the secondary calls across all the frequencies enables the network to place most secondary devices in channels where they can coexist with a cellular handset that is not in the blocking area, which enhances the internal frequency reuse. However, as the number of secondary calls increases, there will be more overlapping of blocking areas where multiple frequencies are unavailable to the cellular calls with the LB scheme. This will increase the blocking probability of cellular calls. A scheme that considers the locations of secondary calls and not just the number of secondary calls underway could perform even better. These superior admission control algorithms are an area for further study.

4.3 Effects of Inaccuracy in Location Measurements

Fig. 2 compares the case where the standard deviation of displacement errors is 100 meters with the case where there is no error. The accuracy level that we choose for comparison is close to the one that the FCC requires for commercial cellular carriers [14]. Even better positioning accuracy is expected in the near future [19]. The distances between a secondary transmitter and a secondary receiver are 100 m and 500 m when the transmit powers of secondary devices (P_{ts}) are 100 mW and 500 mW, respectively. Other parameters are the

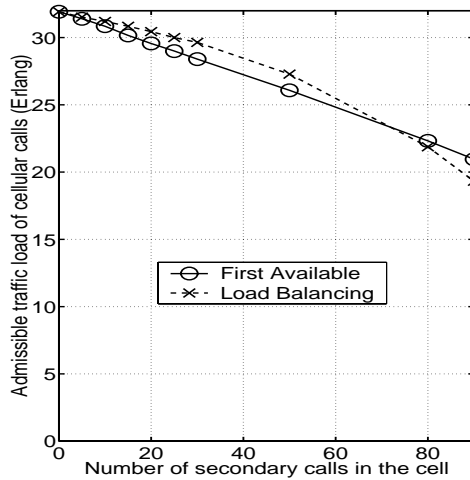


Figure 1: Effect of secondary calls on capacity for cellular customers, with both FA and LB frequency assignment schemes.

same as those in Fig. 1. We only show results with the LB scheme in Fig. 2 because it is better than the FA scheme in this case [18].

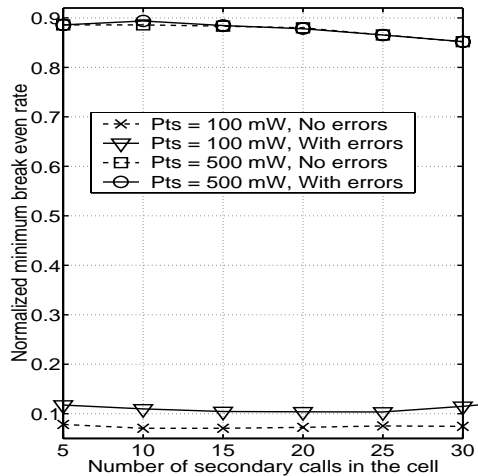


Figure 2: A comparison between two levels of accuracy in location measurement (LB scheme)

It can be observed from this figure that with the accuracy level that current positioning technology can achieve, our scheme still performs well. In addition, in the case where the transmit power of a secondary device is high, the impact of positioning inaccuracy is negligible because the minimum separation required between a cellular handset and a secondary transmitter becomes very large compared to the level of inaccuracy.

4.4 Economic Implications

In Fig. 3, we consider the scenario where a carrier offers secondary access service that can provide an ideal bit rate of 1 Mbps and a transmission distance of 500 m. Each secondary device transmits at 100 mW, and uses 200 kHz of bandwidth, which we have shown to be effective values in this scenario [18]. We only show the FA scheme because it is better than the LB scheme at this particular transmission distance and transmit power [18].

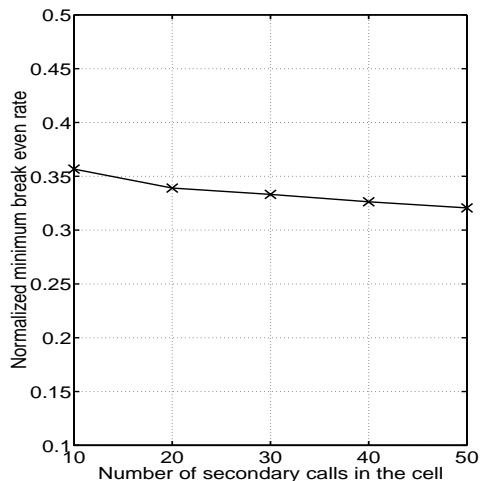


Figure 3: Normalized minimum break even rate when transmission distance is 500 m, and each device transmit at 100 mW. (FA scheme).

We can observe that it is inexpensive to offer a secondary access service. The amount that a carrier needs to charge each secondary call is only about one third of the rate it would normally charge a cellular call. Moreover, the curve looks relatively flat, which suggests that a carrier does not have to get many secondary customers before it can break even.

In Fig. 4, we use the secondary transmit power that minimizes the expected blocking probability of cellular calls for each transmission distance. There are 10 secondary calls underway in the cell. Other parameters are the same as those of Fig. 3. Again, we only show the FA scheme because it is better in all the transmission distances and transmit powers considered in this scenario [18]. We see that a carrier can offer a secondary access service inexpensively at any transmission distance. Even if the transmission distance is as high as 1 km, the carrier only needs to charge about two thirds the per-minute rate it would normally charge a cellular call. Moreover, the normalized minimum break even rate increases roughly linearly as the transmission distance increases, even though we have assumed an omnidirectional antenna, so the area affected by the secondary transmitter increases with distance squared.

In Fig. 5, we choose the combination of SIR and bandwidth that yields the desired bit rate and minimizes the expected call blocking probability of cellular calls. Other parameters are the same as those of Fig. 3. It can be observed that the normalized minimum break




Figure 4: Normalized minimum break even rate at different transmission distances when each device transmits at optimal power

even rate increases linearly when the bit rate is between 1 Mbps and 2 Mbps. However, it looks flat when the bit rate is above 3 Mbps. This is because the system is at the point where the required SIR of each secondary call is so high that each channel of GSM can only accommodate one secondary call and no primary calls can share the channel. Hence, there is no interference from the secondary user's perspective.




Figure 5: Normalized minimum break even rate at different bit rate

5 Payment Mechanism

License-holders will obviously need payment from secondary users to justify the added complication, and the slight loss in revenue from primary users. If software-defined radios allow secondary users to jump from band to band, secondary markets will be most efficient if any carrier can accept payment from any secondary user and any secondary user can make payment to any carrier. Thus, some kind of payment system standard would be helpful. The payment system should have the following properties.

- It should not be necessary for a secondary user to communicate directly with an agent of the payment system. In many cases, the secondary user will have no communications links available other than the signaling channel to the license-holder, which in this case means a short wireless link to the cellular base station.
- To prevent secondary users from fraudulently obtaining services, it must be possible to determine remotely that the payment is good, perhaps by authenticating the secondary user using public key encryption, or perhaps by authenticating the payment itself. The latter approach where the payment rather than the payer is scrutinized is appropriate for a “digital cash” system, where a bit string has value because it has been “minted” (or digitally signed) by a bank [20, 21]. In such a system, the payee requires no knowledge of the payer, which is an advantage in this context.
- Transactions should generate records that cannot be altered without detection, perhaps using digital signatures, so any disputes can be resolved, and fraud can be detected.

Payment systems have emerged in recent years with these properties. A good example is Pay Cash [21], which allows its users to exchange digital currency. It uses public key encryption to verify that the digital currency is valid and digital signatures to create tamper-proof records.

6 Conclusion

Licensing spectrum gives license-holders some exclusivity, thereby guaranteeing quality of service, but exclusivity also leads to inefficient use of spectrum. Unlicensed bands promote sharing, but do not provide adequate quality of service for some applications. Each of these regulatory paradigms have their place, but for applications that require sporadic access to spectrum and for which quality of service guarantees are important, licensed spectrum with real-time secondary markets may be the best solution.

In this scheme, primary and secondary spectrum users share spectrum, thereby increasing spectral efficiency and alleviating spectrum scarcity. Unlike today’s bands that allow secondary access, primary and secondary users would coordinate directly, making it possible to protect quality of service for both primary and secondary users. In this explicit

coordination, the license-holder runs an *admission control* algorithm, which only allows secondary users access to spectrum when quality of service of both primary and secondary would be adequate. The license-holder also uses an intelligent *frequency assignment algorithm* for determining the frequency at which a secondary user should be allowed to operate. Coordination also gives the license-holder an opportunity to charge a fee, which provides incentive to maximize spectrum utilization. Moreover, secondary users dynamically request access to spectrum when and only when spectrum is needed, whether it is needed for one second or one hour.

If secondary devices are equipped with GPS receivers, the admission control and frequency assignment algorithms can allow a great deal of frequency reuse without excessive interference. If secondary devices also employ software-defined radios to jump from band to band, these devices can seek out a band that is currently free, significantly increasing the likelihood that spectrum will be available to secondary devices when needed with negligible impact on license-holders.

For the special case where the license-holder is a GSM-based cellular carrier, we have demonstrated that secondary users can have little impact on a license-holder, so a license-holder can allow many secondary users to operate while reducing the number of cellular customers only slightly, and without reducing quality of service. These results were achieved with two simple frequency assignment algorithms, so it is clear that even better results are possible. 100 meters of inaccuracy in the positioning technology does not significantly undermine the performance of this scheme. Moreover, accuracy in these devices is likely to improve in the coming years [19].

We further show that the cellular carrier profits from this arrangement even if the price for secondary access is quite low. Moreover, the minimum break-even price is relatively insensitive to the number of secondary users, so it is not necessary to get a large number of secondary customers before this approach becomes economically viable.

It should be noted that the only cost considered in our analysis is the opportunity cost of using spectrum for secondary users instead of primary users. There would of course be other costs, ranging from additional capital investment in equipment to marketing of this new service, and the economic viability of secondary access may change if these costs are large.

Our results also indicate that the price of secondary access should be roughly proportional to bandwidth, and roughly proportional to distance from secondary transmitter to receiver. The latter is somewhat surprising because our model used omnidirectional antennas, so the area covered is proportional to distance squared.

Finally, we addressed the challenge of transferring funds from secondary user to license-holder. We describe some of the requirements of a settlement mechanism for this environment, where there may be multiple license-holders and secondary users with different technologies, and where secondary devices have no direct connection to an infrastructure that would allow payment. We presented one system called Pay Cash [21] that can overcome these challenges with a digital currency.

Carriers benefit from the ability to offer a new revenue-generating service. Secondary users obviously benefit, because they get access to spectrum that would otherwise be unavailable with a quality guarantee. Even cellular customers would probably benefit in the end, since added demand from any kind of customer will cause the carrier to expand capacity, which tends to drive cost per customer down. However, this requires analysis which is beyond the scope of this paper. If everyone could benefit, real-time secondary markets deserve serious consideration.

Before opening the doors for real-time secondary markets of this kind, regulators must make sure that the license-holder can be held accountable if a secondary user fails to follow regulatory constraints. For example, what happens if a secondary device transmits at a greater power than is allowed, and who is liable, and how can the offending party be identified? It is not practical to notify the regulator every time a secondary user requests access to spectrum, as might be implied in [3]. Perhaps it is sufficient to obtain a permit from the regulator to be a secondary user in certain bands at any time(s) in the coming year. The authentication mechanisms built into the payment system could be used to verify that this user really does have a permit. Perhaps this process can be bypassed entirely when a secondary user has equipment that has already been approved by the regulator to operate as a secondary user in a given set of frequency bands. This issue also deserves further investigation.

Regulators must also determine exactly what rights can be granted to secondary users, at least in bands where the license-holder does not have complete flexibility. For example, if a license-holder is restricted to FM broadcasting, can it allow secondary users to establish temporary point-to-point microwave links in its frequency allocation? Doing so could change the interference landscape for systems in adjacent frequency bands, systems in nearby regions, or systems that might someday share the same frequency band and geographic region with an FM broadcaster through traditional secondary access. Before allowing real-time secondary markets, the regulator must either grant the license-holder complete flexibility, which in effect grants more spectrum resources, or the regulator must carefully devise technical constraints on the behavior of secondary users in the band. This is also an area that deserves further study.

References

- [1] Federal Communications Commission, “Secondary Markets Spectrum Initiative Policy Statement,” Policy Statement (FCC 00-401).
- [2] Federal Communications Commission, “Promoting Efficient Use of Spectrum Through Elimination of Barriers to the Development of Secondary Markets,” Notice of Proposed Rule Making (FCC 00-402).
- [3] Federal Communications Commission, “FCC Adopts Spectrum Leasing Rules and Streamlined Processing for License Transfer and Assignment Applications, and Pro-

poses Further Steps to Increase Access to Spectrum Through Secondary Markets,” News Release, May 2003.

- [4] J. M. Peha, “Developing Equipment and Services for Shared Spectrum: Is It A Good Gamble?,” *International Engineering Consortium Annual Review of Communications*, vol. 50, 1997, pp. 955-58.
- [5] M. Sirbu, H. Salgado, and J. M. Peha, “A Narrow Band Approach to Efficient PCS Spectrum Sharing Through Decentralized DCA Access Policies,” *IEEE Personal Communications*, vol. 4, no. 1, Feb. 1997, pp. 24-34, www.ece.cmu.edu/~peha/papers.html
- [6] D. P. Satapathy and J. M. Peha, “Performance of Unlicensed Devices with a Spectrum Etiquette,” *Proc. IEEE GLOBECOM*, Nov. 1997, pp. 1645-49, www.ece.cmu.edu/~peha/papers.html
- [7] D. P. Satapathy and J. M. Peha, “Etiquette Modifications For Unlicensed Spectrum: Approach and Impact,” *Proc. IEEE Vehicular Technology Conference (VTC-98)*, May 1998, pp. 272-6, www.ece.cmu.edu/~peha/papers.html
- [8] D. P. Satapathy and J. M. Peha, “A Novel Co-existence Algorithm for Unlicensed Fixed Power Devices,” *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, 2000, pp. 1403-09, www.ece.cmu.edu/~peha/papers.html
- [9] D. P. Satapathy and J. M. Peha, “A Novel Co-existence Algorithm for Unlicensed Variable Power Devices,” *Proc. IEEE International Conference on Communications (ICC)*, vol. 9, 2001, pp. 2845-49, www.ece.cmu.edu/~peha/papers.html
- [10] D. P. Satapathy and J. M. Peha, “Dangers of Excessive Resource Consumption in Unlicensed Spectrum,” *ACM Wireless Networks*, in press.
- [11] J. M. Peha, “Spectrum Management Policy Options,” *IEEE Communications Surveys*, vol. 1, no.1, Fourth Quarter 1998, www.ece.cmu.edu/~peha/papers.html
- [12] J. M. Peha, “Wireless Communications and Coexistence for Smart Environments,” *IEEE Personal Communications*, vol. 7, no.7, Oct. 2000, pp. 66-68, www.ece.cmu.edu/~peha/papers.html
- [13] Shared Spectrum Company, “Comments of Shared Spectrum Company,” Comments in response to FCC Docket No. 00-47, June 2000.
- [14] Federal Communications Commission, “FCC Wireless 911 Requirements,” www.fcc.gov/911/enhanced.
- [15] Federal Communications Commission, “E911 Phase II Decisions,” www.fcc.gov/911/enhanced.

- [16] M. I. Silventoinen and T. Rantalainen, "Mobile Station Emergency Locating in GSM," *Proc. IEEE International Conference on Personal Wireless Communications*, 1996, pp. 232-8.
- [17] H. Laitinen, J. Lahteenmaki, and T. Nordstrom, "Database Correlation Method for GSM Location," *Proc. Vehicular Technology Conference*, vol. 4, 2001, pp. 2504-8.
- [18] S. Panichpapiboon, "Providing Secondary Access to Licensed Spectrum Through Coordination," Master's thesis, Carnegie Mellon University, May 2002.
- [19] E. A. Bretz, "X marks the spot, maybe," *IEEE Spectrum*, April 2000, pp. 26-36.
- [20] D. Chaum, A. Fiat, and N. Naor, "Transaction Systems To Make Big Brother Obsolete," *Communications of the ACM*, Vol. 28, No. 5, Oct. 1985, pp. 1030-44.
- [21] J. M. Peha and I. M. Khamitov, "Pay Cash: A Secure Efficient Internet Payment System," *Proc. Fifth International Conference on Electronic Commerce*, Oct. 2003, www.ece.cmu.edu/~peha/papers.html