Available online at www.sciencedirect.com

SCIENCE DIRECT®

Telecommunications Policy 28 (2004) 603–618

TELECOMMUNICATIONS POLICY

www.elsevierbusinessandmanagement.com/telpol

Real-time secondary markets for spectrum

Jon M. Peha*^,1^, Sooksan Panichpapiboon[2]

Department of Electrical and Computer Engineering, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA

Received 1 October 2003; received in revised form 1 May 2004

Abstract

Spectrum licensing enables quality of service guarantees, but often leads to inefficient use of spectrum. Unlicensed spectrum promotes efficiency through sharing, but quality of service cannot be guaranteed, which is a serious problem for some applications. Such applications may be better served by a *real-time secondary market*, where secondary users ask the license-holder for temporary access to spectrum as needed. Access is granted when and only when quality of service requirements can be met for both license-holder and secondary users. This paper quantitatively assesses real-time secondary markets for the special case of a cellular license-holder. It demonstrates that many secondary users can access spectrum with little impact on the primary cellular customers, and that cellular carriers profit even if the price for secondary access is quite low. Finally, it addresses the challenges of transferring funds from secondary user to license-holder.

Keywords: Spectrum licensing; Secondary market; Unlicensed spectrum; Shared spectrum; Underlay; Spectrum scarcity

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 is the legacy of spectrum regulations designed for old technology. Regulatory reform could unleash new

---

*Corresponding author. Tel.: +1-412-268-7126.

E-mail addresses: peha@cmu.edu (J.M. Peha), sooksan@andrew.cmu.edu (S. Panichpapiboon).

[1] Professor of Electrical Engineering and Public Policy at Carnegie Mellon University (CMU), Associate Director of CMU Center for Wireless and Broadband Networks, www.ece.cmu.edu/~peha.

[2] Graduate Student at Carnegie Mellon University, Department of Electrical and Computer Engineering, www.andrew.cmu.edu/~sooksan.

0308-5961/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.telpol.2004.05.003

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) (2000a, b, 2003). 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. This 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 (Peha, 1997; Sirbu, Salgado, & Peha, 1997). 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 methods to address this problem, including an effective etiquette (Satapathy & Peha, 1997, 1998, 2000, 2001, in press). Giving users the ability to deploy and move transmitters anywhere at any time without first obtaining 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 (Peha, 1998). 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 provided by DSL), so they may not be happy with access over unlicensed spectrum. For applications like these, this paper proposes 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 other 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 provides license-holders with an incentive to share the spectrum, and it provides secondary users with an 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 obtain 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 FCC has lately been considering the possibility of secondary markets where license-holders negotiate with secondary users (FCC, 2000a, b). Very recently (FCC, 2003), they announced that some forms of secondary markets will be allowed. The details are still forthcoming. At minimum, this is an important step toward the vision described in this paper. The first innovators are likely to lease spectrum for months or years. However, these proceedings may not anticipate a real-time market as discussed in this paper, where spectrum is leased for hours, minutes, or even seconds. The announcement defined a lease to be "short term" if it is under 360 days, and a license-holder is expected to notify the FCC 21 days before a secondary user begins operation.

Any license-holder could consider making spectrum available in a real-time secondary market, regardless of the reason for which the license-holder uses that spectrum. 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 (Peha, 1998, 2000). 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. This paper considers 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. In off-peak hours when few primary users are active, it is obvious that granting access to secondary users makes sense, but what about peak hours? The paper shows 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 locations. 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 obtain the desired quality of service guaranteed in one frequency band, the device can jump to another band (Shared Spectrum Company, 2000).

New Internet payment systems also constitute an enabling technology by providing an easy method 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 the proposed scheme. Discussion of the precise conditions for call admission takes place 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, the paper concludes 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, the paper examines the case where the primary user is a GSM-based cellular network operator. A basic assumption is that the cellular carrier is capable of locating the position of each primary device, which is consistent with the FCC's E911 requirement (FCC, 1996). A number of techniques for locating a mobile device in a GSM network has been proposed (Laitinen, Lahteenmaki, & Nordstrom, 2001; Silventoinen & Rantalainen, 1996). The paper restricts itself 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.

Two assumptions are made for secondary devices. First, secondary users only use downstream channels of the cellular network. This is also a conservative assumption, effectively making only half of the carrier's spectrum available for sharing. Secondly, 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 devices 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. Analysis commences here 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 Panichpapiboon (2002) 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. The assumption is 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. It 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, such that the SIR of the downlink meets a threshold level X. Assuming that the 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 the 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. Since X is the SIR required by a primary device, it means that the ratio between the power a primary device receives from the base station and the interference power must be less than X. If a secondary device transmits with power P_{ts}, the interference power received by a primary device in this case is P_{ts}/d_{min}^n. It follows that:

$$\frac{\text{signal}}{\text{interference}} = \frac{P_{max}/C^n}{P_{ts}/d_{min}^n} = X. \tag{1}$$

Solving (1), d_{min} can be written as

$$d_{min} = C\left(\frac{XP_{ts}}{P_{max}}\right)^{1/n}. \tag{2}$$

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

Let Z be the tolerable interference of a secondary device and let w be the distance between a secondary device and a primary base station, then it is required that $P_{tb}/w^n \leqslant Z$, where P_{tb} is the power transmitted by the primary base station. Since power control is used, P_{tb} will depend on the distance between a primary device and the base station. As this distance increases, the power transmitted by the base station increases. Let d_{max} be the maximum distance a primary handset can be from the base station, such that SIR on the secondary link meets a threshold level Z. The transmit power of the base station in this case is $P_{tb} = \left(P_{max}/C^n\right)d_{max}^n$. Consequently, the required condition can be expressed as

$$\frac{\left(P_{max}/C^n\right)d_{max}^n}{w^n} = Z. \tag{3}$$

Solving (3), d_{max} can be written as

$$d_{max} = wC\left(\frac{Z}{P_{\max}}\right)^{1/n}. \tag{4}$$

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. As an example, see Fig. 1. 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

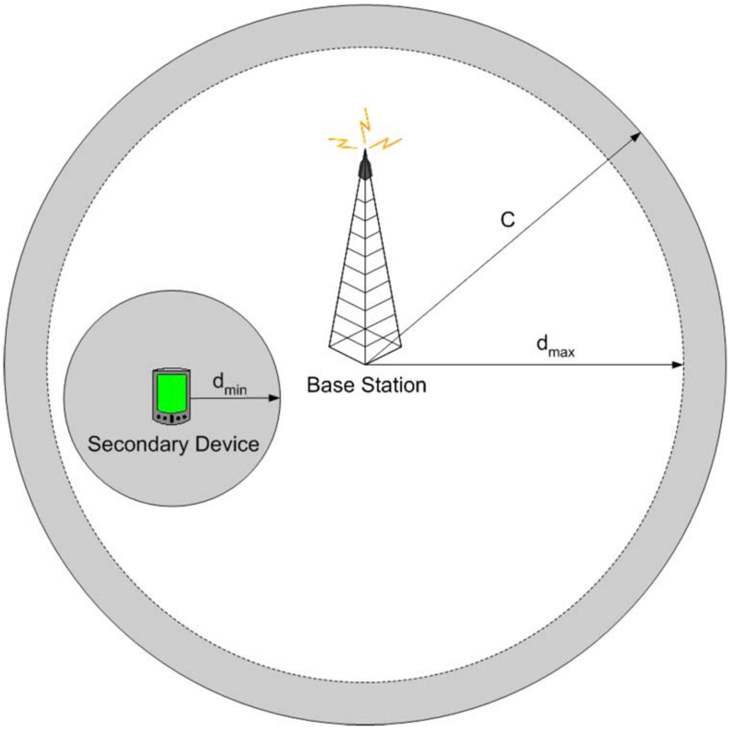

Fig. 1. Blocking area is the shaded region.

- 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 Panichpapiboon, 2002). (Note that generally $E_{05}\leqslant 0$ and $E_{95}\geqslant 0$.)

4. Impact of secondary users on quality of service

This section quantifies the impact of allowing secondary devices to use the spectrum licensed to the cellular network. It evaluates the extent to which carriers would be forced to serve fewer cellular customers, and whether carriers would have a financial incentive to try this form of spectrum sharing.

The assumption is that the system would operate at full capacity even without secondary users, i.e. cellular customers are already experiencing the maximum tolerable blocking probability. (Otherwise, the impact of secondary users could be negligible.) Cellular quality of service (blocking probability) is not allowed to be degraded, so admitting secondary users means reducing the number of primary customers. This reduction will be quantified.

Another assumption is that secondary devices are stationary during periods when they are using spectrum. In addition, it is assumed for this analysis that the 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 in this paper. Blocking probability of cellular calls is calculated as if primary devices are fixed during the calls. Call arrivals of the primaries are assumed to follow a Poisson process, call holding time is exponentially distributed, and primary calls are uniformly distributed throughout the cell. A further assumption is 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 here the assumption is that the location distributions of these transmitters are uniform and independent. Results are achieved through simulation. (See Panichpapiboon (2002) 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, the following values are used unless stated otherwise, all of which are reasonable for a GSM network. The maximum transmit power of a base station is 30 W. The cell radius is 2 km. The 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, system performance is shown 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

Three kinds of metrics are used in measuring how the 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.

The expected call blocking probability is calculated 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 from a uniform distribution. 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 ±5% of the value shown.

Admissible traffic load is the amount of traffic that a network can carry while maintaining a constant quality of service for cellular customers, i.e. a constant expected call blocking probability. The slope of the graph of admissible load from primary users versus the number of secondary users is an indicator of the opportunity cost of allocating spectrum to secondary users; this slope is essential in the calculation of optimal pricing (Wang, Peha, & Sirbu, 1997), although optimal pricing is outside the scope of this paper. 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 ±5% with a 95% confidence interval.

If no secondary calls are allowed, a cellular carrier operating at peak capacity with maximum blocking probability B will have a revenue per minute of $PA_0(1-B)$, where P is a per-minute rate that a carrier normally charges for a cellular call, and A_0 is the admissible traffic load of cellular calls when no secondary calls are present. (Blocked calls produce no revenue.) If N secondary calls are underway, then revenue from cellular calls falls to $PA(1-B)$, where A is admissible traffic load from cellular calls. However, each of the N secondary calls also yield revenue. Thus, for a cellular carrier to gain the same revenue as in the case where there is no spectrum sharing, the following equations must hold:

$$PA_0(1-B) = PA(1-B) + SN, \tag{5}$$

$$\frac{S}{P} = \frac{(1-B)(A_0-A)}{N}, \tag{6}$$

where S is the minimum break-even rate. The ratio in (6) is referred to as the “normalized minimum break-even rate”.

Note that P is the expected revenue derived per minute from a single cellular user in a busy hour, as this is the revenue lost when the number of cellular minutes is reduced. In service plans where revenue is a complex combination of per-minute charges and flat monthly fees, and where some users want peak hour and some off-peak service, the revenue per minute in peak hour is more difficult to calculate.

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. Two simple schemes are introduced: “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, and assigning the first frequency encountered where interference is within the acceptable limits.

It can be observed from Fig. 2 that when there are few secondary calls in the cell, the LB scheme performs better than the FA scheme. This is because in a situation when there are few secondary calls, spreading items 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. 3 compares the case where the standard deviation of displacement errors is 100 m with the case where there is no error. The accuracy level that is chosen for comparison is close to the one that the FCC requires for commercial cellular carriers (FCC, 1996). Even better positioning accuracy is expected in the near future (Bretz, 2000). The distances between a secondary transmitter and a secondary receiver are 100 and 500 m when the transmit powers of secondary devices (P_{ts}) are 100 and 500 mW, respectively. Other parameters are the same as those in Fig. 2. Only results with the LB scheme are shown in Fig. 3 because it is better than the FA scheme in this case (Panichpapiboon, 2002).

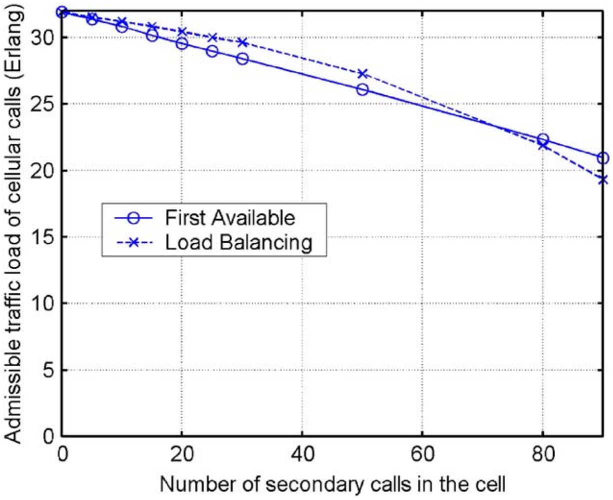

Fig. 2. Effect of secondary calls on capacity of cellular customers, with both FA and LB frequency assignment schemes.

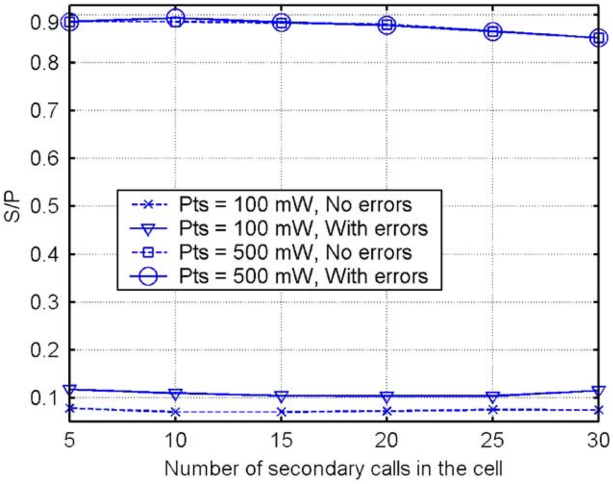

Fig. 3. A comparison between two levels of accuracy in location measurement (LB scheme).

It can be observed from this figure that with the accuracy level that current positioning technology can already achieve, the scheme performs well, averting concern about GPS inaccuracy. 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. 4, 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 is considered. Each secondary device transmits at 100 mW, and uses 200 kHz of bandwidth, which has been shown to be effective values in this scenario (Panichpapiboon, 2002). The FA scheme is shown because it is better than the LB scheme at this particular transmission distance and transmit power (Panichpapiboon, 2002).

It can be observed that it is possible to offer inexpensive secondary access service and still increase revenue. This is due to the fact that, in most cases, primary and secondary devices can coexist in the same block of spectrum; therefore, a carrier gains additional revenue from secondary access while serving a primary user. 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. 5, the secondary transmit power that minimizes the expected blocking probability of cellular calls for each transmission distance is used. There are 10 secondary calls underway in the cell. Other parameters are the same as those in Fig. 4. Again, only the FA scheme is shown because it is better in all the transmission distances and transmit powers considered in this scenario (Panichpapiboon, 2002).

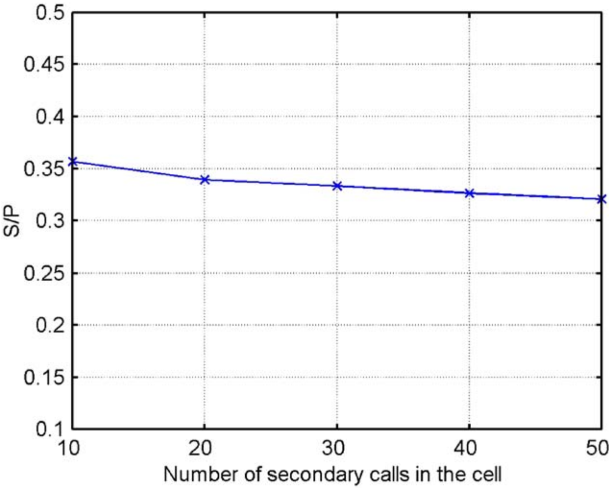

Fig. 4. Normalized minimum break-even rate when transmission distance is 500 m, and each device transmits at 100 mW (FA scheme).

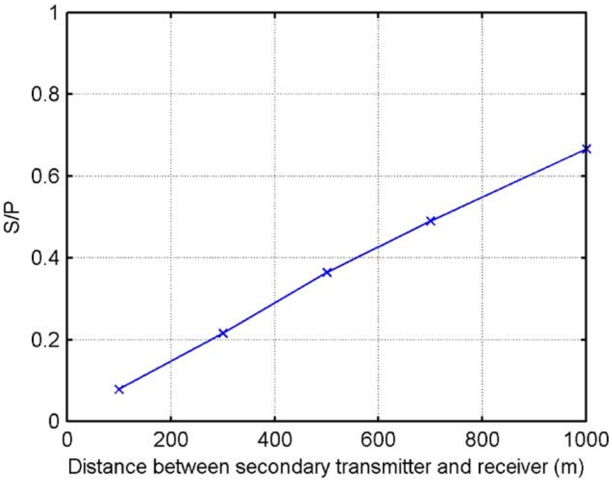

Fig. 5. Normalized minimum break-even rate at different transmission distances when each device transmits at optimal power.

A carrier can profitably offer an inexpensive secondary access service 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 with the assumption of an omnidirectional antenna, so the area affected by the secondary transmitter increases with distance squared.

In Fig. 6, the combination of SIR and bandwidth is chosen that yields the desired bit rate and minimizes the expected call blocking probability of cellular calls. Other parameters are the same as

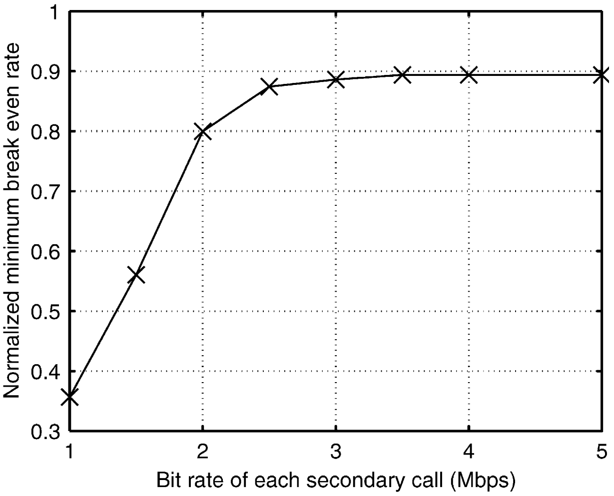

Fig. 6. Normalized minimum break-even rate at different bit rate.

those of Fig. 4. It can be observed that the normalized minimum break-even rate increases linearly when the bit rate is between 1 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 GSM channel 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.

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 (Chaum, Fiat, & Naor, 1985; Peha & Khamitov, 2003). 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. In particular, digital cash systems are well suited to meeting the first two conditions, where the payer has no direct connection with the payment system but all payments are validated. There is one such system called Pay Cash (Peha & Khamitov, 2003) that offers the strongest protection against disputes through its tamper-proof records. This is accomplished through a unique approach that inextricably links a piece of electronic cash to the associated transaction record. Without such a link, a service-provider might claim that a particular service was never paid for, and that any payment made by that customer was actually for a previous service. Conversely, a customer can falsely claim that another payment was really for this service.

The link between cash and transaction record is achieved as follows (Peha & Khamitov, 2003). Electronic cash systems typically associate serial numbers with electronic cash, so the system can block attempts to spend the same cash twice. In Pay Cash, this serial number is also a public key, and the corresponding secret key is used to sign the agreement between payer and payee which will eventually become part of the transaction record.

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 incentives to maximize spectrum utilization. Moreover, secondary users dynamically request access to spectrum when and only when spectrum is needed, whether it is needed for 1 s or 1 h.

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. In effect, by increasing frequency reuse, this technology increases the economy of scope of offering these two different services on the same infrastructure. 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, the paper has quantitatively 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 and sharing only in with downstream cellular links, so it is clear that even better results are possible. In all, 100 m 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 coming years (Bretz, 2000).

The paper further shows 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 such customers before this approach becomes economically viable.

It should be noted that the only cost considered in the present analysis is the opportunity cost of using spectrum for secondary instead of primary users. There would of course be other costs, ranging from additional capital investment in equipment to marketing of this new service.

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

Finally, the paper addressed the challenge of transferring funds from secondary user to license-holder. It describes 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. The paper presented one system called Pay Cash (Peha & Khamitov, 2003) 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 clearly 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. It could be sufficient to obtain a permit from the regulator to be a secondary user in a number of bands and 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 considerably more spectrum resources (without an auction or the usual consideration), 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—on a band-by-band basis.

References

Bretz, E. A. (2000). X marks the spot, maybe. *IEEE Spectrum*, *37*(4), 26–36.

Chaum, D., Fiat, A., & Naor, N. (1985). Transaction systems to make big brother obsolete. *Communications of the ACM*, *28*(5), 1030–1044.

Federal Communications Commission (FCC). (1996). *FCC adopts rules to implement enhanced 911 for wireless services*. FCC News Release.

Federal Communications Commission (FCC). (2000a). *Secondary markets spectrum initiative policy statement*. Policy Statement (FCC 00-401), November 9.

Federal Communications Commission (FCC). (2000b). *Promoting efficient use of spectrum through elimination of barriers to the development of secondary markets*. Notice of Proposed Rule Making (FCC 00-402), November 9.

Federal Communications Commission (FCC). (2003). *FCC adopts spectrum leasing rules and streamlined processing for license transfer and assignment applications, and proposes further steps to increase access to spectrum through secondary markets*. FCC News Release.

Laitinen, H., Lahteenmaki J., & Nordstrom, T. (2001). Database correlation method for GSM location. *Proceedings of the vehicular technology conference* (pp. 2504–2508).

Panichpapiboon, S. (2002). *Providing secondary access to licensed spectrum through coordination*. Master's thesis, Carnegie Mellon University, 2002.

Peha, J. M. (1997). Developing equipment and services for shared spectrum: Is it a good gamble? *Proceedings of the international engineering consortium annual review of communications* (pp. 955–958).

Peha, J. M. (1998). Spectrum management policy options. *IEEE Communications Surveys*, *1*(1), 2–8.

Peha, J. M. (2000). Wireless communications and coexistence for smart environments. *IEEE Personal Communications*, *7*(7), 66–68.

Peha, J. M., & Khamitov, I. M. (2003). Pay cash: A secure efficient Internet payment system. *Proceedings of the fifth international conference on electronic commerce*.

Satapathy, D. P., & Peha, J. M. (1997). Performance of unlicensed devices with a spectrum etiquette. *Proceedings of the IEEE GLOBECOM* (pp. 1645–1649).

Satapathy, D. P., & Peha, J. M. (1998). Etiquette modifications for unlicensed spectrum: Approach and impact. *Proceedings of the IEEE vehicular technology conference* (pp. 272–276).

Satapathy, D. P., & Peha, J. M. (2000). A novel co-existence algorithm for unlicensed fixed power devices. *Proceedings of the IEEE wireless communications and networking conference* (pp. 1403–1409).

Satapathy, D. P., & Peha, J. M. (2001). A novel co-existence algorithm for unlicensed variable power devices. *Proceedings of the IEEE international conference on communications* (pp. 2845–2849).

Satapathy, D. P., & Peha, J. M. (in press). Dangers of excessive resource consumption in unlicensed spectrum. *ACM Wireless Networks*.

Shared Spectrum Company. (2000). Comments of Shared Spectrum Company. Response to FCC Docket No. 00-47.

Silventoinen, M. I., & Rantalainen, T. (1996). Mobile station emergency locating in GSM. *Proceedings of the IEEE international conference on personal wireless communications* (pp. 232–238).

Sirbu, M., Salgado, H., & Peha, J. M. (1997). A narrow band approach to efficient PCS spectrum sharing through decentralized DCA access policies. *IEEE Personal Communications*, *4*(1), 24–34.

Wang, Q., Peha, J. M., & Sirbu, M. A. (1997). Optimal pricing for integrated-services networks. In J. Bailey, & K. McKnight (Eds.), *Internet economics* (pp. 353–376). Cambridge, MA: MIT Press.