Limits on Communications in a Cognitive Radio Channel

Natasha Devroye, Patrick Mitran, and Vahid Tarokh, Harvard University

ABSTRACT

In this article we review FCC secondary markets initiatives and how smart wireless devices could be used to increase spectral efficiency. We survey the current proposals for cognitive radio deployment, and present a new, potentially more spectrally efficient model for a wireless channel employing cognitive radios; the *cognitive radio channel*. This channel models the simplest scenario in which a cognitive radio could be used and consists of a 2 Tx, 2 Rx wireless channel in which one transmitter knows the message of the other. We obtain fundamental limits on the communication possible over such a channel, and discuss future engineering and regulatory issues.

INTRODUCTION

Recent FCC measurements have indicated that 90 percent of the time, many licensed frequency bands remain unused [1]. As user demands for data services and data rates steadily increase, efficient spectrum usage is becoming a critical issue.

In order to better utilize the licensed spectrum, the FCC has recently launched a Secondary Markets Initiative [2], whose goal is to "remove regulatory barriers and facilitate the development of secondary markets in spectrum usage rights among Wireless Radio Services." This proposal introduces the concept of "dynamic spectrum licensing," which implicitly requires the use of cognitive radios to improve spectral efficiency. *Cognitive radio*, a term first coined by Mitola [3], is a low-cost, highly flexible alternative to the classic single-frequency-band single-protocol wireless device. By sensing and adapting to its environment, a cognitive radio is able to cleverly avoid interference and fill voids in the wireless spectrum, dramatically increasing spectral efficiency.

Although the gains to be made by the combination of cognitive radios and *secondary spectrum licensing* seem intuitive, the *fundamental* theoretical limits of the gains to be made by this coupling have only recently been explored [4, 5]. This motivates the writing of this article, where we review the basics of cognitive radio and the FCC initiatives they opportunistically exploit. Furthermore, the current state of the art on the theoretical limits of wireless channels employing cognitive radios are laid out, as well as a novel idea for an achievable rate region that more fully exploits the capabilities of cognitive radios. In short, the question of how much data can be reliably transmitted over the newly defined *cognitive radio channel* is posed in information theoretic terms, in order to conclusively explore the limits of this new channel. This channel is modeled as a two-sender, two-receiver interference channel, with one twist: the *genie*. Suppose a (possibly non-cognitive) radio is transmitting. A cognitive radio that wishes to transmit may listen to the wireless channel, and can obtain the signal of the currently transmitting user. The genie idealizes message knowledge, and noncausally gives the incumbent cognitive radio full, noncausal knowledge of the existing transmitters' messages. We argue why this is a viable model to explore and what conclusions may be drawn from these results. Approaching the problem from an information theoretic angle is novel, as the limited research on cognitive radios tends to come from a more practical protocol-oriented perspective. We finally explore some of the regulatory and engineering aspects that must be addressed in order to realize these gains.

COGNITIVE RADIO: THE SMART APPROACH

Over the past few years, the incorporation of software into radio systems has become increasingly common. This has allowed for faster upgrades, and has given these wireless communication devices more flexibility, and the ability to transmit and receive using a variety of protocols and modulation schemes (enabled by reconfigurable software rather than hardware). Furthermore, as the name suggests, such radios can even become "cognitive" and, as dictated by the software, adapt their behavior to their wireless surroundings without user intervention. According to the FCC, software defined radio (SDR) encompasses any "radio that includes a transmitter in which operating parameters such as frequency range, modulation type or maximum output power can be altered by software without making any changes to hardware components that affect the radio frequency emissions." Mitola [3] took the definition of an SDR one step further, and envisioned a radio that could make decisions as to the network, modulation, and/or coding parameters based on its surround-

This material is based upon research supported by the National Science Foundation under the Alan T. Waterman Award, Grant no. CCR-0139398.

 0163-6804/06/$20.00 © 2006 IEEE

ings, and called such a "smart" radio a cognitive radio. Such radios could even make decisions based on the availability of nearby collaborative nodes, or on the regulations dictated by their current location and spectral conditions.

One of the main players in the early development of software defined radios was the U.S. Department of Defense Joint Tactical Radio System (JTRS) Program. The JTRS developed a software architecture known as the Software Communications Architecture (SCA), into which different hardware components may be integrated. SCA was later adopted by commercial industry through a non-profit international organization aimed at promoting SDR technology, called the SDR Forum. In an alternative and parallel approach, the open source GNU radio project hopes to encourage research and development of SDRs, allowing anyone to contribute their own code to the already existing openly available software. In the EU, the End-to-End Reconfigurability (E^2R) Project [6] aims at realizing the full benefits of the diversity within the radio eco-space, composed of a wide range of systems such as cellular, fixed, wireless local area, and broadcast. The systems they intend to develop will "provide common platforms and associated execution environments for multiple air interfaces, protocols, and applications, which will yield to scalable and reconfigurable infrastructure that optimize resource usage through the use of cognition-based methods." Other SDR research efforts include the collaboration of Tektronix with Virginia Tech's Mobile and Portable Radio Research Group, as well as a new National Science Foundation "Research in Networking Technology and Systems" (NeTS) program.

Cognitive radio technology is perfectly suited to opportunistically employ the wireless spectrum. Their *frequency agility*, *dynamic frequency selection*, *adaptive modulation*, *transmit power control*, *location awareness*, and *negotiated use* — meaning ability to incorporate agreements into their behavior — all allow for very flexible spectrum use. In essence, cognitive radios could skillfully navigate their way through interference, and greatly improve spectral efficiency. The FCC, very enthusiastic about these possibilities, is now vigorously altering their regulations to allow for more flexible use of the licensed wireless spectrum.

Secondary Markets: Encouraging Efficiency

Since 2000, the FCC has actively been developing a Secondary Markets Initiative, as well as various rulemaking releases regarding the use of cognitive radio technologies. They are interested in removing unnecessary regulatory barriers to new secondary-market-oriented policies such as:

- **Spectrum leasing**: Allowing unlicensed users to lease any part or all of the spectrum of a licensed user.
- **Dynamic spectrum leasing**: Temporary and opportunistic usage of spectrum rather than a longer-term sublease.
- **"Private commons"**: A licensee could allow unlicensed users access to his/her spectrum without a contract, optionally with an access fee.

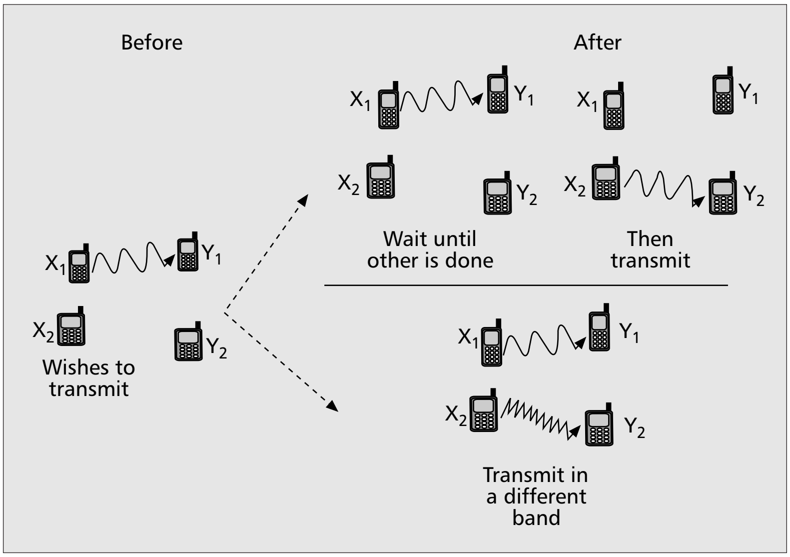

Figure 1. *Current proposals for dynamic spectrum leasing involve two schemes: when a cognitive radio X_2 wishes to transmit to Y_2 and a possibly non-cognitive X_1 is already transmitting to Y_1, it can either wait until X_1 has completed its transmission (time division, as in the top figure), or possibly transmit at a different frequency band (frequency division, as in the bottom figure). In either case time or frequency division is employed, rather than sharing the time/frequency spectrum.*

- **Interruptible spectrum leasing**: would be suitable for a lessor that wants a high level of assurance that any spectrum temporarily in use, or leased, to an incumbent cognitive radio could be efficiently reclaimed if needed. A prime example would be the leasing of the generally unoccupied spectrum allotted to the U.S. government or local enforcement agencies, which in times of emergency could be quickly reclaimed. Interruptible spectrum leasing methods resemble those of *spectrum pooling*. Reference [7] provides a nice overview of spectrum pooling and solutions to some of the associated technical aspects.

In current FCC proposals on opportunistic channel usage, the cognitive radio listens to the wireless channel and determines, in either time or frequency, which resources are unused [1]. It then adapts its signal to fill this void in the spectrum domain, by transmitting either at a different time or in a different band, as shown in Fig 1. Thus, a device transmits over a certain time or frequency band *only when no other user does*. Another potentially more flexible, general, and spectrally efficient approach would be to allow *two users to simultaneously transmit* over the same time or frequency. Under this scheme, a cognitive radio will listen to the channel and, if sensed idle, could proceed as in the current proposals (i.e., transmit during the voids). On the other hand, if another sender is sensed, the radio may decide to proceed with simultaneous transmission. The cognitive radio need not wait for an idle channel to start transmission. Some questions that arise with this new model are: is this spectrally more efficient than time sharing the spectrum? What are the achievable rates at which two users could transmit, and how does this compare to when the devices are not cognitive radios, yet still proceed in the same fashion?

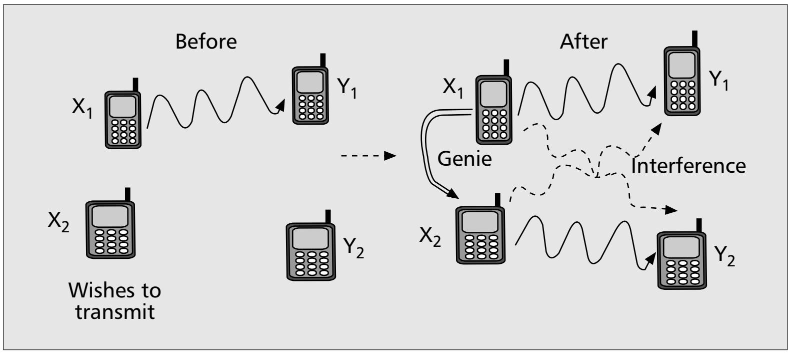

Figure 2. *The cognitive radio channel is defined as a two-sender* (X_1, X_2), *two-receiver* (Y_1, Y_2) *interference channel in which the cognitive radio transmitter* X_2 *is noncausally given by a genie the message* X_1 *plans to transmit.* X_2 *can then either mitigate the interference it will see, aid* X_1 *in transmitting its message, or, as we propose, a smooth mixture of both.*

What regulatory issues will be faced? What engineering problems will need to be solved for this to enter into the mainstream?

Cognitive Radio Channels: Exploiting Flexible Spectrum Usage

Cognitive radios have the ability to listen to the surrounding wireless channel, make decisions on the fly, and encode using a variety of schemes. In order to fully exploit this, first consider the simplest example, shown in Fig. 2, of a channel in which a cognitive radio device could be used in order to improve spectral efficiency. As shown on the left, suppose sender X_1 is transmitting over the wireless channel to receiver Y_1, and a second incumbent user, X_2, wishes to transmit to a second receiver, Y_2. In current secondary spectrum licensing proposals, the incumbent user X_2, a cognitive radio that is able to sense the presence of other transmitting users, would either wait until X_1 has finished transmitting before proceeding, or possibly transmit over a different frequency band. Rather than forcing X_2 to wait, in [4] we have suggested allowing X_2 to *simultaneously* transmit with user X_1 at the same time in the same band of frequencies. The wireless nature of the channel will make interference between simultaneously transmitting users unavoidable. However, by making use of the capabilities of a cognitive radio, we have shown that the cognitive radio is able to potentially mitigate the interference. The question we pose is thus: what are the fundamental communication limits of such a two-sender, two-receiver scheme in which at least one user, the incumbent transmitter, is a cognitive radio? To more precisely define the problem, as well as to solve it from a theoretical perspective, we translate it into the language of information theory.

Cognitive Radio Channels: Capacity vs. Achievable Regions

One of the many contributions of information theory is the notion of *channel capacity*. Qualitatively, it is the maximum rate at which information may be sent *reliably* over a channel. When there are multiple simultaneous information streams being transmitted, we can speak of capacity regions as the maximum set of all reliable rates that can be simultaneously achieved. For example, the capacity region of the channel depicted in Fig. 2 is a two-dimensional region, or set of rates (R_1, R_2), where R_1 is the rate between ($X_1 \rightarrow Y_1$) and R_2 is the rate between ($X_2 \rightarrow Y_2$). For any point (R_1, R_2) inside the capacity region, R_1 on the *x*-axis corresponds to a rate that can be reliably transmitted simultaneously, over the same channel, with R_2 on the *y*-axis. There exist many channels whose capacity regions are still unknown. For such channels, tight inner and outer bounds on this capacity region are research goals. An inner bound is also called an *achievable rate/region*, and consists of suggesting a particular (often random) coding scheme and proving that the claimed rates can be reliably achieved, that is, that the probability of a decoding error vanishes with increasing block size. Notice that this guarantees the existence of schemes which can reliably communicate at these rates. Random coding does not construct explicit practical schemes, and does not guarantee that better schemes do not exist. We will demonstrate our achievable region [4, 5] for the two-sender 2 receiver case in which at least one sender is a cognitive radio.

The Genie: Message Knowledge Idealization

What differentiates the cognitive radio channel from a basic 2 sender, 2 receiver interference channel is the *message knowledge* of one of the transmitters. This message knowledge is possible due to the properties of cognitive radios. If X_2 is a cognitive radio, and is geographically close to X_1 (relative to Y_1), then the wireless channel ($X_1 \rightarrow X_2$) could be of much higher capacity than the channel ($X_1 \rightarrow Y_1$). Thus, in a fraction of the transmission time, X_2 could listen to, and obtain the message transmitted by X_1. It could then employ this message knowledge — which translates into exact knowledge of the interference it will encounter — to intelligently try to mitigate it. Although we have used transmitter proximity to motivate the message idealization assumption, and have proposed a particular transmission scheme for this scenario, different relative distances between transmitting and receiving nodes could dictate different schemes, as is investigated in [8]. Important to note is that our scheme is beneficial mostly in the weak interference case, as the strong [9, 10] and very strong [11] interference channels have known capacity regions and known ways of achieving them. The relative node positions will determine what type of interference channel results.

We introduce the genie so as to idealize the message knowledge of sender X_2. That is, we suppose that rather than causally obtaining the message X_1 is transmitting, a fictitious genie hands X_2 this message. Notice that X_1 is not given X_2's message, so we have an asymmetric problem. This idealization will provide an upper bound to any real-world scenario, and the solutions to this problem may provide valuable insight to the fundamental techniques that could be employed in such a scenario. We also expect that under suitable proximity of the two transmitters, this bound is nearly achievable. The techniques used in obtaining the limits on communication for the channel employing a genie could be extended to provide achievable regions for the case in which X_2 obtains X_1's message causally. We have suggested causal schemes in [12].

Achievable Region of the Cognitive Radio Channel

A cognitive radio channel [4] is a two-transmitter, two-receiver classical information theoretic interference channel in which sender 2 (a cognitive radio) obtains, or is given by a genie, the message sender 1 plans to transmit. The scenario is illustrated in Fig. 2. The cognitive radio may then simultaneously transmit over the same channel, as opposed to waiting for an idle channel as in a traditional cognitive radio channel protocol. Although the capacity region of the formulated cognitive radio channel at first glance seems to be a simple problem, it is also still an open one. Thus, an intuitively pleasing achievable region for the rates (R_1, R_2) at which X_1 can transmit to Y_1, and X_2 to Y_2, simultaneously, was constructed in our previous work in [4] and improved in [5]. This construction merges ideas used in dirty-paper (or Gel'fand-Pinsker) coding [13] with the Han and Kobayashi achievable region construction [9] for the interference channel, as well as the relay channel. When X_2 has *a-priori* knowledge of what X_1 will transmit, or the interference it will encounter, one can think of two possible courses of action:

1 **Selfishly** try and mitigate the interference. This can be done using a dirty paper coding technique [13]. In this case, X_2 is layering on his own independent information to be transmitted to Y_2. This strategy yields points of higher R_2 and lower R_1 in the cognitive channel region of Fig. 3.

2 **Selflessly** act as a relay to reinforce the signal of user X_1. Such a scheme, although it does not allow X_2 to transmit its own independent information, seems intuitively correct from a fairness perspective. That is, since X_2 infringes on X_1's spectrum, it seems only fair that X_1 should somehow benefit. This strategy yields points of high R_1 and lower R_2 in the cognitive channel region of Fig. 3.

In [5] we demonstrate an achievable region that smoothly interpolates between these two schemes. The resulting achievable region in the presence of additive white Gaussian noise is plotted as the "cognitive channel region" in Fig. 3. There, we see four regions. The time-sharing region (1) displays the result of pure time sharing of the wireless channel between users X_1 and X_2. Points in this region are obtained by letting X_1 transmit for a fraction of the time, during which X_2 refrains, and vice versa. These points would be amenable to the current proposals on secondary spectrum licensing. The interference channel region (2) corresponds to the best known achievable region of the classical information theoretic interference channel. In this region, both senders encode independently, and there is no message knowledge by either transmitter. The cognitive channel region (3) is the achievable region proposed in our prior work [5] and described here. In this case X_2 received the message of X_1 non-causally from a genie, and X_2 uses a coding scheme which combines interference mitigation with relaying the message of X_1. As expected, the region is convex and smooth. One can think of the convexity as a consequence of time sharing: if any two (or more) schemes achieve certain rates, then by time-sharing these schemes, any convex combination of the rates can be achieved. The region is smooth since our scheme actually involves power sharing at the coding level, which tends to yield rounder edges. We see that both users, not only the incumbent X_2 which has the extra message knowledge, benefit from using this scheme. This is as expected, as the selfish strategy boosts R_2 rates, while the selfless one boosts R_1 rates; thus, gracefully combining the two will yield benefits to both users. The presence of the incumbent cognitive radio X_2 can be beneficial to X_1, a point which is of practical significance. This could provide yet another incentive for the introduction of such schemes. The modified MIMO bound region (4) is an outer bound on the capacity of this channel: the 2×2 multiple-input multiple-output (MIMO) Gaussian broadcast channel capacity region, where we have restricted the form of the transmit covariance matrix to be of the form

$$\begin{pmatrix} P_1 & c \\ c & P_2 \end{pmatrix}$$

to more closely resemble our constraints, intersected with the capacity bound on R_2 for the channel for $X_2 \rightarrow Y_2$ in the absence of interference from X_1.

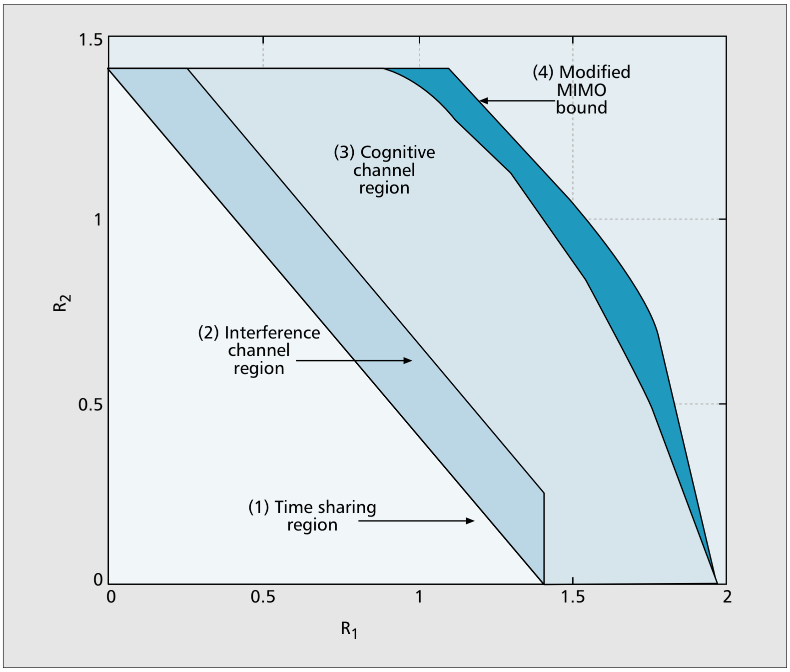

Figure 3. *Rate regions* (R$_1$, R$_2$) *for different two-sender, two-receiver wireless channels. Region (1) is the time sharing region of two independent senders. Region (2) is the best known achievable region for the interference channel, as calculated by Han and Kobayashi [9]. Region (3) is the achievable region described here and in [5] for the cognitive radio channel. Region (4) is an outer bound on the cognitive radio channel capacity. All simulations are in AWGN, with sender powers 6 and noise powers 1. The crossover parameters in the interference channel are 0.55 and 0.55.*

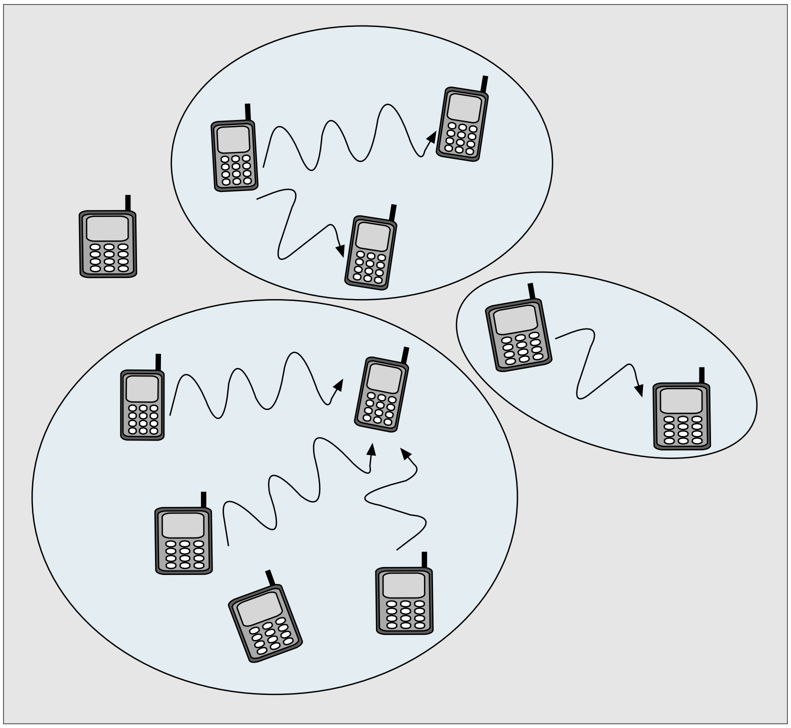

■ **Figure 4.** *A wireless network at a given instance in time can be decomposed into non-interfering groups, as shown here, and further divided into clusters, as done in [16].*

Extensions: Cognitive Radio Networks

The simple 2×2 wireless channel employing and exploiting cognitive radios can be generalized to larger *cognitive networks*, where we abstract the asymmetric form of transmitter cooperation to a general type of *cognitive behavior*. In our work [16], when given three pieces of information about a wireless network — the information graph (which indicates independent information streams), the interference graph (which indicates which transmission can be heard by whom) ,and information about which nodes are cognitive — we derive a *cognitive graph*. This graph consists of a set of non-interfering groups, as shown in Fig. 4, which can be further decomposed into possible overlapping clusters. Different levels of transmitter cooperation within and between clusters can be investigated. The *inter/intra-cluster competitive*, *cooperative*, and *cognitive* behavior in wireless networks, as shown in Fig. 5, are defined in [16]. These represent three types of transmitter cooperation and encompass a wide range of classical information theoretic channels. We define *intercluster cognitive behavior* as simultaneous transmission of messages by two or more clusters in which some clusters know (given by a genie) the messages to be transmitted by other clusters. Similarly, *intracluster cognative* behavior is when nodes within one cluster obtain the messages of other nodes within that same cluster and simultaneously transmit. An achievable region for the inter-cluster behavior of two multiple access channels is constructed in the authors' prior work [5].

Practical Considerations and Future Directions

As demonstrated by Fig. 3, when two users must share a channel, there is an incentive for both the cognitive and non-cognitive users of a wireless channel to employ a scheme other than time sharing. Better rates can be achieved by both users. However, for such a cognitive scheme to become reality, many practical engineering aspects must be overcome. First, an efficient coding scheme that combines a dirty-paper-like method with a relay-like technique will have to be constructed. Practical coding schemes for channels with known side information at the transmitter have recently received a great deal of attention [14]. Such methods could potentially be modified to the current needs. The achievable region calculated requires full channel knowledge, an idealistic assumption. The construction of good codes that perform well even if partial or noisy channel state information is available is another hurdle to overcome. In addition, the genie idealization must be removed. In our extended work on cognitive radio channels [12] we provide two-phase protocols (listening phase, cognitive transmission phase) for which cognitive user X_2 may causally obtain user X_1's message. Although this is a start, alternate causal protocols will need to be developed. As the genie represents an idealization, causal schemes may use the genie-aided achievable region as an outer bound. Theoretical bounds on what can be achieved in the causal case still remains an open question. Another interesting engineering aspect would be to see the intuitive trade-off between (partial) message knowledge and achievable rates. Then a cognitive radio could decide when it has obtained a sufficient portion of the message (or with sufficient reliability) to operate at the desired point in the region. The transmission scheme derived here assumed asymmetry in the capability of the two transmitting devices. However, in a future in which all devices are smart, new possibilities for simultaneous transmission arise. Should both exchange messages then transmit, or should transmission occur in a cognitive fashion? Do better schemes become possible when the problem becomes symmetric?

Once the basic two-sender, two-receiver case is solved from a practical perspective, scaling this behavior to large networks of cognitive radios must also be considered. Given a general network with cognitive nodes, we must determine which, how many, and how nodes should best collaborate to transmit their respective messages. Finding protocols that perform and scale well under both cognitive radio capabilities and regulatory constraints will be of vital importance.

Regulatory Considerations

The FCC is trying to make the process of secondary spectrum licensing as painless as possible. In addition, they are aggressively working on encouraging the development and use of cognitive radio technology. The proposed secondary spectrum licensing, some of which lies in the VHF and UHF television bands, has caused some controversy [15], and the FCC is welcoming comments on issues relating to secondary licensing of spectrum. In the EU, the E^{2}R [6] project is also considering the associated regulatory issues. The FCC envisions at least four possible scenarios in which cognitive radio technologies could be used to improve spectral efficiency [1]. First, a licensee would use cognitive radios internally to increase efficiency within its own spectrum. Second, they could be used in easing secondary spectrum licensing, between a licensee and a third party. Third, they could facilitate automated frequency coordination among licensees of a co-primary license. Finally, in the situation mostly considered here, a cognitive radio could act as an unlicensed device opportunistically employing the spectrum in time. Our proposal would require clarification of this final use: rather than restrict cognitive radios to time sharing the channel, they must obtain the right to *concurrent* spectrum use, a more delicate regulatory question. Since choice of the modulation and coding parameters would allow operation anywhere inside the (R_1, R_2) achievable rate region, measures must be taken to ensure that the incumbent cognitive radio, which will have permission to simultaneously transmit, will not abuse this right and adversely affect current users.

The FCC is also currently investigating what kind of technology in cognitive radios could guarantee the immediate release of any borrowed spectrum for *interruptible spectrum* leasing or *spectrum pooling* [7]. This is particularly relevant in the context of governmental emergency bands, which for the most part remain unused and would be prime candidates for secondary licensing or dynamic spectrum sharing. Such agencies will be reluctant to proceed with secondary licensing unless such a guarantee can be made. These issues have been addressed in [1], and could be extended to controlling incumbent cognitive radios in other scenarios as well.

Conclusion

In this article we review the basics of cognitive radios and recent FCC secondary spectrum licensing initiatives for increasing the spectral efficiency of wireless channels. We propose an alternate scheme for exploiting both the cognitive radio capabilities and the new, more flexible licensing agreements. This motivates the definition of the cognitive radio channel, a two-transmit, two-receive interference channel in which one user knows the message to be transmitted by the other. Fundamental limits on communication are established for such channels, and engineering and regulatory aspects in order to approach these limits are discussed.

References

[1] FCC Spectrum Policy Task Force, "FCC Report of the Spectrum Efficiency Working Group," Nov. 2002, http://www.fcc.gov/sptf/files/SEWGFinalReport1.pdf

[2] Secondary Markets Initiative, http://wireless.fcc.gov/licensing/secondarymarkets/1

[3] J. Mitola III, "Cognitive Radio: An Integrated Agent Architecture for Software Defined Radio," Ph.D. Thesis, KTH Royal Inst. Technology, Stockholm, Sweden, 2000.

[4] N. Devroye, P. Mitran, and V. Tarokh, "Achievable Rates in Cognitive Radio Channels," *39th Annual Conf. Info. Sci. and Sys.*, Mar. 2005.

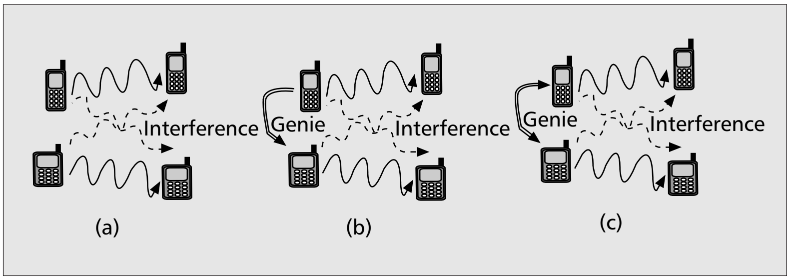

Figure 5. *Wireless clusters of nodes can behave in 3 fashions: a) they can compete for the wireless resources (competitive); b) partially cooperate (cognitive); c) fully cooperate during transmission (cooperative). Here intercluster, or transmitter behavior between clusters, is demonstrated.*

[5] N. Devroye, P. Mitran, and V. Tarokh, "Cognitive Multiple Access Networks," *Proc. IEEE Int'l. Symp. Info. Theory*, Sept. 2005.

[6] End-to-End Reconfigurabiity, http://phase2.e2r.motlabs.com/

[7] T. A. Weiss and F. K. Jondral, "Spectrum Pooling: An Innovative Strategy for the Enhancement of Spectrum Efficiency," *IEEE Commun. Mag.*, Radio Commun. Supp., Mar. 2004, pp. S8–S14.

[8] C. T. K. Ng and A. Goldsmith, "Capacity Gain from Transmitter and Receiver Cooperation," *Proc. IEEE Int'l. Symp. Info. Theory*, Sept. 2005.

[9] T. S. Han and K. Kobayashi, "A New Achievable Rate Region for the Interference Channel," *IEEE Trans. Info. Theory*, vol. IT-27, no.1, Jan. 1981, pp. 49–60.

[10] H. Sato, "The Capacity of Gaussian Interference Channel Under Strong Interference, " *IEEE Trans. Info. Theory*, vol. IT-27, no. 6, Nov. 1981.

[11] A. B. Carleial, "Interference Channels," *IEEE Trans. Info. Theory*, vol. IT-24, Jan., 1978, pp. 60–70.

[12] N. Devroye, P. Mitran, and V. Tarokh, "Achievable Rates in Cognitive Radio Channels," *IEEE Trans. Info. Theory*, vol. 52, no. 5, May 2006.

[13] M. H. M. Costa, "Writing on Dirty Paper," *IEEE Trans. Info. Theory*, vol. IT-29, May 1983, pp. 439–41.

[14] R. Zamir, S. Shamai, and U. Erez, "Nested Linear/Lattice Codes For Structured Multi-Terminal Binning," *IEEE Trans. Info. Theory*, vol. 48, no. 6, June 2002.

[15] M. J. Marcus, "Unlicensed Cognitive Sharing of the TV Spectrum: the Controversy at the Federal Communications Commission," *IEEE Commun. Mag.*, May 2005.

[16] N. Devroye, P. Mitran, and V. Tarokh, "On Cognitive Graphs: Decomposing Wireless Networks," to appear, *Proc. Crowncom*, Mykonos, Greece, June 2006.

Biographies

NATASHA DEVROYE [S'01] (ndevroye@deas.harvard.edu) received a Bachelor's degree in electrical engineering in 2001 from McGill University, Montreal, Quebec, Canada. She received her Master's degree in 2003 from the Division of Engineering and Applied Sciences at Harvard University, Cambridge, Massachusetts, where she is currently working toward her Ph.D degree. In 2005 she interned as a research scientist for Intel Corporation in the Radio Communications Laboratory. Her research interests include multi-user information theory, wireless networks, as well as cognitive radio communications.

PATRICK MITRAN [S'01] (mitran@deas.harvard.edu) received Bachelor's and Master's degrees in electrical engineering in 2001 and 2002, respectively, from McGill University. He is currently working toward a Ph.D. degree in the Division of Engineering and Applied Sciences, Harvard University. In 2005 he interned as a research scientist for Intel Corporation in the Radio Communications Laboratory. His research interests include iterative decoding, distributed/joint source-channel coding, cooperative and cognitive communications, as well as information theory.

VAHID TAROKH [M'95, SM'03] (vahid@deas.harvard.edu) is a Perkins Professor of Applied Mathematics and a Vinton Hayes Senior Research Fellow of Electrical Engineering at Harvard University, where he defines and supervises research in the general areas of signal processing, communications, and networking.