Performance analysis of a multi-antenna framework for spectrum reuse

Omar Mohammed Bakr
Mark Christopher Johnson
Ben Wild
Kannan Ramchandran

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2008-101
http://www.eecs.berkeley.edu/Pubs/TechRpts/2008/EECS-2008-101.html

August 20, 2008

Performance analysis of a multi-antenna framework for spectrum reuse

Omar Bakr Mark Johnson Ben Wild Kannan Ramchandran

Electrical Engineering and Computer Sciences, UC Berkeley
Email:{ombakr, mjohnson, benwild, kannanr}@eecs.berkeley.edu

***Abstract*— This paper proposes a new framework for spectrum reuse. Existing architectures have centered on secondary users (cognitive radios) that can reliably sense primary users and opportunistically transmit, without directly interacting with the primary system. We argue that a paradigm in which the primary and secondary systems *cooperate* can result in reduced interference to primary users and more predictable access for secondary users. Because this architecture gives the primary system full control over spectrum sharing, it could be more favorable in the current economic and political environment.**

We illustrate a concrete instance of our framework by showing how secondary radios can reuse the entire uplink channel of a cellular network. We also demonstrate a computationally efficient beamforming algorithm, which enables the coexistence of the two systems. The proposed architecture requires only modest changes to the primary infrastructure, and is shown to achieve an interference rejection of up to 20 dB in most practical scenarios.

I. Introduction

The current system of spectrum allocation has resulted in the vast under-utilization of frequency resources [1], [2]. While all bands below 3 GHz have been allocated [3], measurements of spectrum usage reveal significant spatial and temporal variations, including large “white spaces” of unused spectrum [4], [5]. In order to more effectively utilize scarce frequency resources, the FCC has issued a Notice of Proposed Rule Making [6], advancing Cognitive Radio (CR) technology as a candidate to implement negotiated or opportunistic spectrum sharing. These cognitive radios would be designed to operate in multiple frequency bands and dynamically adapt their transmission to their environment. Building practical cognitive radio systems is a significant technical challenge.

A system architecture designed to enable spectrum reuse must satisfy several requirements in order to be commercially viable. Most important, the amount of interference and service degradation experienced by the primary (legacy) system as a consequence of the presence of the secondary (cognitive radio) system must be kept below a tolerable level. It is crucial that the primary system have the ability to *control* and minimize the interference that it experiences. Also, the secondary system must be assured of consistent and *predictable* access to the spectrum, in order to provide a meaningful quality of service (QoS) to its users. Finally, deploying the system must be economically feasible. The cost of the new hardware required by both the primary and secondary systems must be acceptably small.

Most spectrum sharing research to date has focused on systems that opportunistically reuse frequency bands by detecting the presence and absence of primary systems [7], [8]. To prevent interference with the primary system, the secondary system must be able to accurately sense the presence of primary users, and then either suppress transmission or control its radiation pattern via beamforming when primary users are present. In this paradigm, the primary is not modified in any way when the secondary system is deployed, and in fact has no knowledge of the presence of secondary users.

There are a number of challenges associated with opportunistic spectrum sharing schemes. There are fundamental limitations to the detection of signals in low SNR environments [9]. Shadowing and deep channel fades will further degrade detector performance. Also, in many primary systems, for example television, the receivers are passive devices. Thus, the secondary will not receive any feedback, even implicitly, from the primary users. Monitoring interference created by the secondary system then requires explicit feedback from the network operator of the primary system, which contradicts the principle of opportunistic reuse. In addition, due to the time-varying nature of wireless channels and the fact that primary users enter and leave the system, the secondary must sense the spectrum continuously, which consumes significant resources.

Due to these factors, opportunistic spectrum sharing based on sensing is often either technically difficult or economically infeasible. While collaboration among secondary users in both sensing [10] and beamforming [11] can improve performance and reduce the requirements on a single radio, such techniques cannot completely eliminate, or even effectively control, interference levels.

In this paper we present a new paradigm for spectrum sharing, based on collaboration between the primary and secondary systems. We argue that a system architecture based on cooperation could exploit the *spatial domain* more effectively than an opportunistic architecture, and provide a higher QoS level to both systems. The cooperative framework provides the primary system with an economic incentive to enable spectrum reuse. Secondary systems, by aggregating spectrum from multiple primary networks, have the capability to provide services that a single primary system could not offer.

In Section II, we present a general framework for spectrum reuse that can be applied to a wide variety of systems. Then, in Section III, we give a detailed treatment of a specific application: the reuse of the uplink channel of a cellular network for both short and long range communication. Section IV describes an efficient beam nulling scheme that uses multiple antennas to minimize the interference to primary users, and gives simulation results of the algorithm. While spectrum sharing is a rapidly evolving field with many open problems, we hope that this work provides the impetus for examining these questions from a new direction and offers some useful guidelines for designing and deploying practical systems.

II. Cooperative framework for spectrum reuse

We propose a new framework for spectrum reuse that is based on cooperation between the primary and secondary systems. In contrast to opportunistic schemes, the primary and secondary collaborate to control interference at primary users and guarantee spectrum access to the secondary system. In general, this cooperation will be implemented via feedback from the primary that informs the secondary when to transmit and how much interference is being generated. The channel used to convey this feedback to the secondary users, as well as the exact nature of the feedback, are application specific. They will depend on several factors, including the size and topology of both systems, the underlying physical layer of the primary, the required bandwidth and QoS for each system, and the degree to which the network operators will invest in hardware upgrades. In our framework, it is beneficial, but not essential, for the secondary users to also be members of the primary system, so called *dual citizens*, to facilitate cooperation and enhance the QoS for secondary users. In this work, we will focus on the dual citizen approach to spectrum reuse.

The fundamental technical tools that enable spectrum sharing in this framework are beamforming and beam nulling. The secondary transmitters will minimize interference to the primary system by adjusting their array patterns so that there are nulls in the directions of the primary receivers. If the secondary users have a sufficient number of antennas, they can further improve spectrum utilization by also beamforming in the direction of their intended receivers.

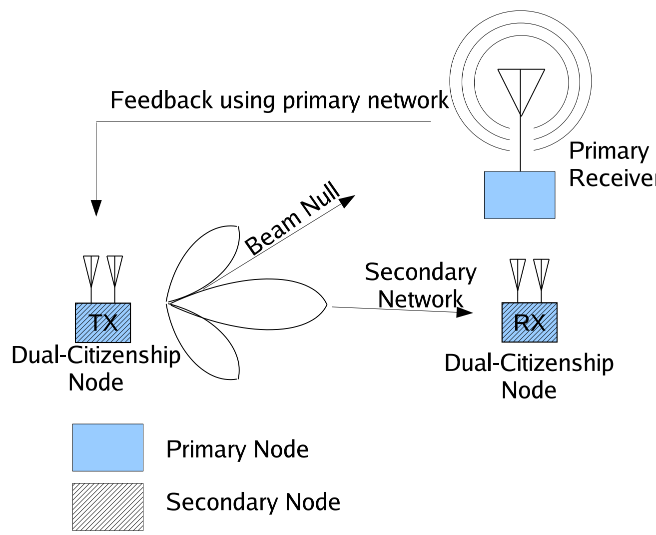

Fig. 1. Basic architecture for cooperative spectrum reuse

The basic elements of this general architecture are summarized in Figure 1. A secondary transmitter (employing multiple antennas) receives feedback from the primary receiver. This feedback assists the secondary in shaping its beam pattern, so that the power radiated in the direction of the primary receiver is effectively zero. In Section III, we demonstrate how this framework can be applied to the specific application of reusing the uplink frequency band of an orthogonal frequency division multiple access (OFDMA) cellular network.

Cooperation has a number of economic, technical, and policy advantages for both primary and secondary systems.

- In this model, there is no need for government regulation. Spectrum sharing is driven purely by economic forces. Without the necessity of complex

and time-consuming government involvement, spectrum reuse systems can be deployed and adapted much more quickly.
- The primary system has *full control* at all times, giving it more protection against service degradations. This in turn gives the primary an economic incentive to accommodate secondary systems.
- Explicit feedback from the primary system enables secondary users to reduce their interference more effectively than in an opportunistic paradigm.
- The additional hardware and complexity required in both systems are quite low.
- Because the scheme is based on existing, mature technologies, it can be deployed more quickly than proposed cognitive radio systems that rely on technical capabilities that have not yet been implemented and tested.

While the benefits of the cooperative framework are clear, a number of challenges also must be overcome before systems based on this paradigm are practical.

- Primary users must be active nodes that have the ability to announce their presence. While this initially appears to be a significant barrier, it may be less of a problem in the future as devices are increasingly connected to the internet via WiFi and bluetooth radios.
- For the secondary system to be practical and economically feasible, primary receivers must be relatively static and geographically sparse, e.g., cellular base stations, satellite base stations, or TV receivers.
- There are also open questions as to how to adapt this model to existing standards, such as WiMAX, that use technologies like MIMO and space-time coding. Spectrum reuse in systems where primary users have multiple antennas could benefit from collaborative and distributed beamforming techniques [11], [12].

In this framework, the secondary system could be either a separate entity or an extension of the primary. For example, the primary network operator could deploy its own secondary system to provide service upgrades to some users, without requiring that all primary users replace their hardware. However, there are other scenarios that might necessitate the deployment of the secondary as an independent system. For example, the operator of the secondary system may be able to use this architecture to aggregate bandwidth from several primary systems and provide services that any single primary is unable to offer[1].

III. Cellular uplink reuse

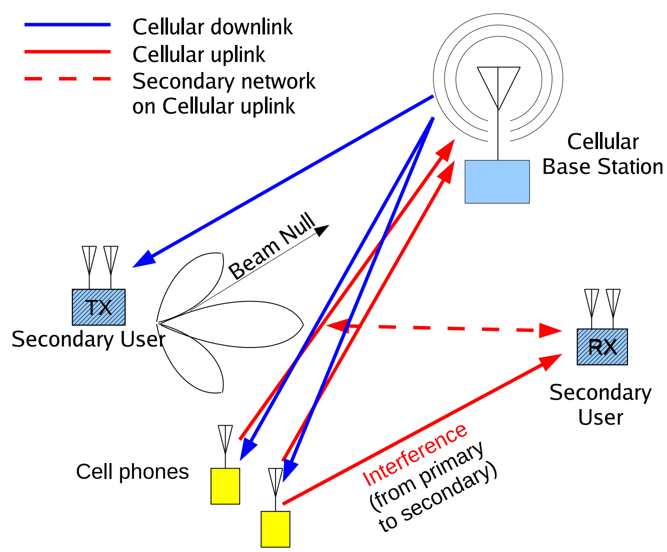

Fig. 2. Cooperative reuse of a cellular uplink channel

The uplink channel of a cellular communication system is a band that can be effectively reused by our proposed framework. This is because on the uplink, unlike the downlink, the primary receivers (base stations) are sparsely distributed and have static locations. In addition, base stations are more easily modified by the network operator than the mobile units.

The basic scenario is depicted in Figure 2. The multi-antenna secondary user connects to the base station as though it were a regular subscriber (a dual citizen), and is allocated channel resources (i.e., time slots, subcarriers, chip sequences, etc.) on both the uplink and the downlink. Once the connection is established, the secondary radio starts transmitting on its allocated channel of the uplink band. The goal of the secondary user is to choose a beamforming weight vector such that the signals from its antenna array cancel out at the base station. Efficient beam nulling schemes are explored in more detail in Section IV. Once the secondary user has chosen a weight vector that results in sufficient signal rejection at the base station, it begins transmitting on the entire uplink band. However, the base station can inform the secondary user to stop transmitting on the entire uplink. Thus, the base station can suppress the secondary system if primary users experience service degradation.

[1]Legal and economic reasons might prevent two primary systems from sharing bandwidth.

In frequency division duplexing (FDD) systems, the secondary user requires explicit feedback from the base station (on the downlink) in order to learn the uplink channel response and compute an appropriate weight vector. On the other hand, in time division duplexing (TDD) systems, the secondary user can learn the channel responses of the uplink by directly estimating the downlink and using channel reciprocity. In the TDD case, the secondary user can transmit on the entire cellular band, but only in uplink time slots.

For a secondary system to be viable, not only does it have to avoid causing interference to the primary system, but it must also be able to suppress interference from the primary transmitters. In the cellular uplink application, the secondary receiver may experience interference caused by some of the primary transmitters (mobile phones) on the uplink. However, there are several reasons why primary interference should not excessively degrade the secondary system. Experimental measurements show that the spectrum usage on the uplink is much less than on the downlink [2]. Also, because the mobile units use a much lower transmit power than the base stations, their impact on the performance of the secondary receivers will be less significant. Furthermore, a secondary radio can use its multiple antennas when receiving to suppress large signal jammers, without the assistance of explicit feedback. Finally, in many cellular standards with universal frequency reuse (e.g. CDMA, OFDMA), the secondary receiver can take advantage of interference averaging [13], which employs coding to reliably communicate when a portion of the time-frequency slots are lost to interference.

A. OFDMA cellular systems

As a concrete example of cellular uplink reuse, we will consider an OFDMA system, which is particularly well suited to this cooperative spectrum reuse framework. In general, higher bandwidth systems require more degrees of freedom to achieve good signal rejection over the entire spectrum[2]. The required complexity can be reduced by using OFDM, which divides a wideband channel into N parallel and orthogonal subchannels, which can be individually nulled. Each subchannel, referred to as a subcarrier, is much smaller than the coherence bandwidth. The subcarriers can be treated as single tap narrowband channels, which can be nulled with fewer antennas than wideband channels.

[2]In scattering environments with high delay spreads, the signal bandwidth can be much larger than the coherence bandwidth. In this case the channel response will be composed of multiple taps. The number of antennas required to null the entire band is proportional to the number of taps.

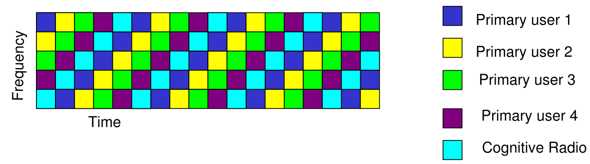

Fig. 3. Multiple access in OFDMA systems

Every subscriber is allocated a time-frequency hopping sequence by the base station on both the uplink and the downlink, as shown in Figure 3. In any given time slot (OFDM symbol), different users transmit on different subcarriers, which implies that the hopping sequences allocated to different subscribers are orthogonal to one another. The hopping sequence is periodic, with period $d \leq N$ where d is usually chosen to be prime [13]. A single period is known as an OFDM block. Each hopping sequence contains d distinct subcarriers, which are usually spread out over the entire band in order to maximize the diversity. To reuse the uplink, the secondary proceeds as follows:

(1) The secondary user requests a hopping pattern from the cellular network and offers the network payment for the use of these resources. The network agrees to the terms and assigns the secondary user a time/frequency hopping pattern. At this point the secondary user is a member of both the primary and secondary systems (a dual citizen).
(2) The secondary radio transmits a pilot sequence on this pattern, and receives feedback from the base station[3].
(3) The secondary radio uses the feedback information to adaptively choose a set of antenna weights that nulls out its signal at the base station. The nulling process is done for each subcarrier individually.
(4) Once the SNR at the base station on all subcarriers falls below a fixed threshold, the secondary radio is permitted to reuse the entire uplink band for a fixed period of time. The secondary can reuse all N subcarriers, since the d subcarriers in its allocated hopping sequence are spread throughout the band. If the level of measured interference is small, the base station can permit the secondary to continue to use the channel.

[3]The exact nature of the feedback depends on the specific beamforming algorithm in use. For example, the feedback might be the signal value at the receiver or the energy in the received signal.

(5) The secondary user continues to transmit a known sequence on its assigned hopping pattern, and receives regular feedback from the base station. The secondary can thus adapt its antenna weights as the channels vary in time.

Because the secondary is nulling each subcarrier individually, and the narrowband subcarriers experience flat fading [13], two antennas are sufficient. Note that in this scenario the secondary users will have to use OFDM to communicate with each other, as the nulling is done on a per subcarrier basis.

In practice, the secondary system can potentially interfere with multiple base stations, not only the base station with which it is registered. This is especially true in high density or sectorized cellular networks. The exact number of base stations depends on the density and network topology, as well as the desired coverage of the secondary system. The scheme described in this section can be used to null multiple base stations, if the secondary radio is assigned the same hopping sequence by each of them. This could be accomplished by having the primary system reserve a fixed sequence at all base stations for use by the secondary system. However, the number of required antennas will also grow linearly with the number of base stations.

Figure 4 summarizes transmit beam nulling in an FDD cellular system[4]. The secondary user transmits to the base station on the hopping sequence of subcarriers it has been assigned[5]. Every subcarrier k at antenna j is premultiplied by a complex weight $c_j[k]$. At each antenna, the subcarriers are transformed into a time domain sequence using the Inverse Discrete Fourier Transform (IDFT), which is then transmitted over the air. The base station receives the time domain signal and transforms it back into the frequency domain using the DFT. The base station feeds back channel information for each subcarrier on the downlink. The feedback information required by the secondary user varies depending on the beamforming or nulling algorithm being used. For example, it could be a quantized version of the channel or simply a single bit indicating changes in SNR [12]. The secondary radio uses this feedback to adapt the amplitudes and the phases of each subcarrier independently.

[4]Some details of the OFDM system that are not relevant to this problem (e.g., cyclic prefix) are omitted.

[5]At the beginning of the process, the secondary user can only transmit on the subcarriers assigned to it by the base station at a given OFDM symbol.

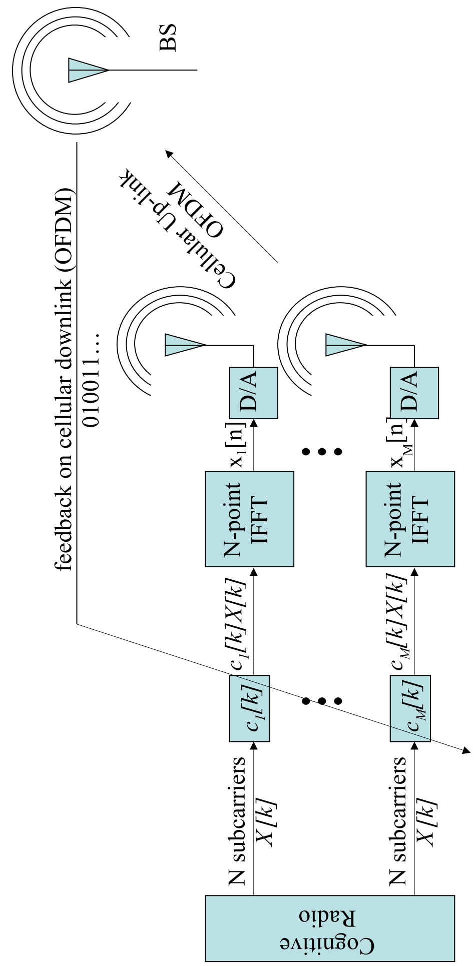

Fig. 4. Beam nulling in OFDMA systems

IV. Beamforming and interference cancellation

To control interference in the primary system, secondary radios exploit multiple antennas and beamforming. The required number of antennas depends on the number of primary receivers with which a secondary radio interferes, which in turn depends on the communication range of the secondary system. For very short range communication, one antenna may be sufficient, and feedback from the primary system allows the secondary to regulate its power. As the desired range of the secondary system increases, however, power control alone cannot prevent interference at the primary receivers. In that case, the secondary radio must have multiple antennas. The interference is minimized by selecting an antenna pattern with nulls in the directions of all primary receivers.

Suitable antenna weights are obtained in a two phase procedure. First, an adaptive filtering algorithm is used to estimate the relevant channel state information. Then these channel states are used to compute the antenna weights.

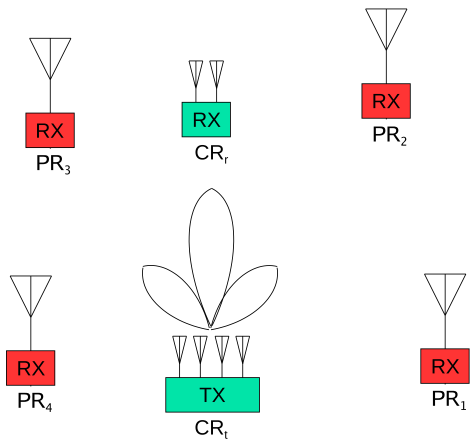

Fig. 5. The cognitive transmitter must null the directions of the four primary receivers.

The general scenario is illustrated in Figure 5. The cognitive radio transmitter CR_t wants to communicate with a cognitive radio receiver CR_r, and can potentially interfere with K primary receivers $PR_1, \ldots, PR_K$. The transmitter has $M > K$ antennas, while all primary receivers have a single antenna. We assume that both systems are narrowband, so that the channels are represented by a single tap in the time domain. The complex channel responses from CR_t to CR_r and PR_i are denoted respectively by

$$\mathbf{h}_r = [h_{r,1} \ldots h_{r,M}]$$

$$\mathbf{h}_i = [h_{i,1} \ldots h_{i,M}] \quad 1 \leq i \leq K$$

In a line of sight AWGN channel, the spatial taps are functions of the angle of arrival and the geometry of the array. On the other hand, in rich scattering environments with large angular spreads, the spatial taps can be modeled by $i.i.d$ complex Gaussian random variables [13].

When CR_t transmits a symbol $x[n]$, it is premultiplied by a complex weight c_j^* at antenna j. The received signal at node PR_i is given by

$$y_i[n] = \mathbf{c}^h \mathbf{h}_i x[n] + \nu[n]$$

where $\mathbf{c} = [c_1 \cdots c_M]$ is the weight vector and $\nu[n]$ is white noise.

A. Optimal beamforming weights

One straightforward method of computing the antenna weights $\mathbf{c}$ is to project $\mathbf{h_r}$ onto the subspace that is orthogonal to the space spanned by the set $\{\mathbf{h_i}, 1 \leq i \leq K\}$, which would perfectly null the directions of all of the primary receivers. However, reducing the interference power well below the noise floor may unnecessarily degrade the desired signal without significantly improving the performance of the primary.

The optimal beamformer is found by choosing weights $\mathbf{c}$ that minimize $E[|y_i[n]|^2]$, while still maintaining an SNR at the cognitive receiver that is sufficient to enable reliable communication. This can be mathematically formalized as minimizing the total interference power $\sum_{i=1}^{K} E[|y_i[n]|^2]$, subject to a constraint on the channel response at the cognitive receiver. We define the $M \times M$ spatial interference plus noise correlation matrix as

$$\mathbf{R_{I+N}} = \sum_{i=1}^{K} \mathbf{h}_i \mathbf{h}_i^h + \sigma_\nu^2 \mathbf{I}$$

where $\mathbf{I}$ is an $M \times M$ identity matrix. The resulting constrained optimization problem is stated as

$$\min \mathbf{c}^h (\mathbf{R_{I+N}}) \mathbf{c} \;\; \text{subject to} \;\; \mathbf{c}^h \mathbf{h}_r = 1$$

The solution to this optimization problem takes the following form [14]:

$$\mathbf{c}_{opt} = \frac{\mathbf{R_{I+N}^{-1}} \mathbf{h}_r}{\mathbf{h}_r^h \mathbf{R_{I+N}^{-1}} \mathbf{h}_r}$$

Note that the denominator is a constant which ensures that the constraint is met. We define the interference rejection of this beamformer toward PR_i (denoted by IR_i) as the ratio of the SNR from CR_t to PR_i with beamforming to the SNR without beamforming (i.e., with $c_j = 1/\sqrt{M}, \;\; 1 \leq j \leq M$). Observing that the magnitude of the channel response with beamforming is $|\mathbf{c}^h \mathbf{h}_i|$ and without beamforming is $|\frac{1}{\sqrt{M}} \sum_{j=1}^{M} h_{i,j}|$, we see that the interference rejection is

$$IR_i = 20 \log \left| \frac{\mathbf{c}_{opt}^h \mathbf{h}_i}{\frac{1}{\sqrt{M}} \sum_{j=1}^{M} h_{i,j}} \right|$$

If the cognitive radio has perfect knowledge of all channel taps, it can use this beamformer to achieve significant interference rejection. In practice, the channel responses must be estimated, and the performance will be degraded due to various factors, such as limited resolution, synchronization errors, and time-variations. However, we will show that in practical scenarios the interference power can be brought below the noise level.

B. Channel estimation via adaptive filtering

The channel responses can be estimated by formulating this problem as an instance of system identification [15] and using an adaptive filter. The system to be identified is a spatial filter, instead of the standard temporal filter. The general identification problem is illustrated in Figure 6. The input $w[n]$ is a white random signal that is generated at the cognitive radio[6]. The unknown system $H(z)$ is the transfer function from the multiple antenna transmitter to the receiver. For example, if we are estimating the channel to PR_i then $H(z)$ is the z-transform of the M tap spatial channel $\mathbf{h}_i$.

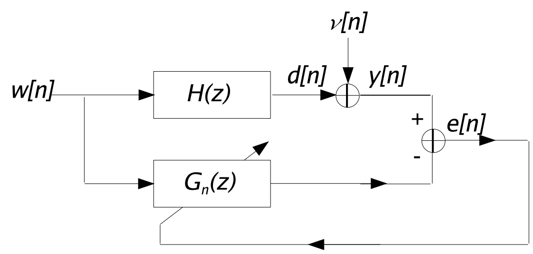

Fig. 6. The transfer function of an unknown system $H(z)$ can be tracked with an adaptive filter $G_n(z)$.

When estimating the channels, CR_t sets the antenna weights to a block of length M of the random sequence $w[n]$, and transmits the pilot $x[n] = 1$. At each time step, the block of $w[n]$ assigned to the weights is shifted to the right by one sample. Thus, the output $y_i[n]$ at PR_i, which is physically produced by M antennas transmitting over single tap channels, mimics the time-domain convolution of the signal $w[n]$ with an M tap filter[7].

The received signal $y_i[n]$ is then fed back to the transmitter through the downlink. The error signal $e[n]$ is used to update the M tap adaptive filter $G_n(z)$. Figure 7 shows how the general system identification paradigm is used to estimate the channels from a multi-antenna transmitter to a receiver. Because $y_i[n]$ is a noisy version of the true desired signal $d[n]$, $G_n(z)$ will not be exactly equal to $H(z)$. However, as our subsequent simulations show, the adaptive filter produces an estimate of sufficiently good quality to enable the cognitive radio to effectively null the primary receivers.

[6] $w[n]$ can be sampled from any distribution (e.g. Gaussian, Bernoulli, etc.).

[7] It is not strictly necessary for the spatial filter to simulate a time domain filter. We can also choose the antenna weights to be $i.i.d.$ in both time and space, and the subsequent algorithm will still robustly estimate the channel.

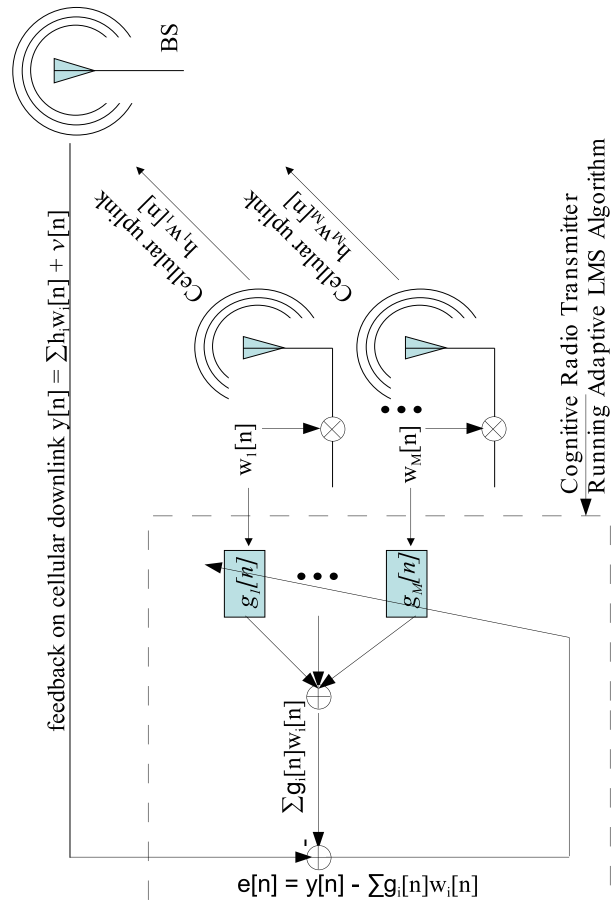

Fig. 7. The cognitive radio uses a random pilot sequence of weights, and employs an adaptive filter to estimate the unknown channel gains.

Note that in Figure 7, $y[n]$ is the signal that is fed back over the downlink. The computation of $e[n]$ and the updating of the adaptive filter taps occurs in the cognitive radio. Therefore, the channels to all primary base stations, and also the cognitive receiver, can be estimated simultaneously, as long as the cognitive radio has been allocated the same channel resources on the uplink by all of the base stations. The cognitive radio simply maintains a separate adaptive filter for each channel that is being estimated. The cognitive radio can also estimate the downlink channels with an adaptive filter, and then use spatial division multiplexing to maximize the feedback rate. Unlike the uplink, the secondary radio can learn the downlink channels without explicit feedback, so long as the primary base stations transmit a known pilot sequence.

In our work, we used the LMS algorithm to estimate the channels, as it is a simple and widely used technique[8]. The convergence time is linear in the number of antennas if the channel taps are $i.i.d.$ While LMS is a mature algorithm whose performance is well understood in terms of convergence rate and error power, we are interested in its performance in the context of beam nulling. In particular, how much interference rejection is sacrificed by using the estimate from the LMS algorithm to compute the antenna weights?

C. Iterative channel estimation

While it is theoretically possible for the cognitive transmitter to concurrently estimate the channels to all primary receivers, as discussed in Section IV-B, implementing such a scheme may not be desirable. In order to accurately estimate the channels, the cognitive radio must be assigned the same channel resources by all primary receivers. In the context of the reuse of the uplink band in an OFDMA cellular network, this means that all of the base stations must assign the cognitive radio the same hopping sequence. This imposes a strong requirement on the primary system, which may not wish to reserve hopping sequences for use by only the secondary system.

On the other hand, if each base station in the OFDMA example assigns a different hopping sequence to the cognitive radio, then the scheme described in Section IV-B suffers from two problems. First, the amount of bandwidth available to the secondary system decreases. When the cognitive radio is forced to reserve a unique hopping sequence for channel estimation for each base station with which it interferes, then the number of hopping sequences available for communication with secondary receivers is reduced. Second, the scheme will result in unacceptable interference levels at the base stations. Imagine that PR_1 assigns a hopping sequence to CR_t, which then begins using that sequence to learn $\mathbf{h_1}$. With high probability, this hopping sequence will interfere with sequences that PR_2 has assigned to its own primary users.

In order to solve these problems in a practical manner, we propose a method of iteratively estimating the channels to the primary receivers. We will consider the scenario in Figure 5, and assume without loss of generality that PR_1 is the primary receiver with the strongest channel to CR_t, PR_2 has the second strongest channel, etc. The cognitive transmitter CR_t initially uses a low transmit power, so that it can only communicate with PR_1. The scheme outlined in Section IV-B is used to learn $\mathbf{h}_1$, the channel weights to PR_1. This will not cause any interference at the other primary receivers, since they are assumed to have weaker channels than PR_1.

Next, the cognitive radio increases its transmit power until it can communicate with two primary receivers. To learn the channels to PR_2, the previous adaptive filtering method is modified. At each iteration of the LMS algorithm, CR_t will generate a white vector $\mathbf{w}$, and then find the component of $\mathbf{w}$ that is orthogonal to $\mathbf{h_1}$. This orthogonal component is then used as the antenna weights by the adaptive filter[9]. In contrast to the algorithm in Section IV-B, the weights used by the LMS algorithm are now a function of the estimates of previous channels. If we express the channel weights to PR_2 as $\mathbf{h_2} = \mathbf{h_2}^{\parallel} + \mathbf{h_2}^{\perp}$, where $\mathbf{h_2}^{\parallel}$ is the component parallel to $\mathbf{h_1}$ and $\mathbf{h_2}^{\perp}$ is the component orthogonal to $\mathbf{h_1}$, then by choosing the training sequence in this manner the adaptive filter will converge to $\mathbf{h_2}^{\perp}$.

The algorithm proceeds in an iterative manner. The cognitive radio increases its power until it can communicate with three primary receivers, and while running the LMS algorithm chooses random weights that are orthogonal to both $\mathbf{h_1}$ and $\mathbf{h_2}^{\perp}$. Thus, CR_t will learn the component of $\mathbf{h_3}$ that is orthogonal to both $\mathbf{h_1}$ and $\mathbf{h_2}$. Our scheme, which learns the orthogonal component of each vector of channel taps, is analogous to the Gram-Schmidt process of computing orthogonal basis vectors.

D. Simulation results

We simulated the performance of the beamformer in Section IV-A in order to evaluate it in realistic environments. We first considered the algorithm in Section IV-B and a time-invariant channel. The plot in Figure 8(a) shows the SNR at one primary receiver after beamforming as a function of the SNR before beamforming. Since the secondary system wants to prevent interference to the primary, the goal of the beamformer is to minimize this output SNR. Different curves represent different numbers of antennas used by CR_t. In all cases there were a total of $K = 4$ primary receivers in the system. Note that the SNR before beamforming is the SNR at which the adaptive filter operates while estimating the channel. Figure 8(b) shows the output SNR when an error is added to the phase and amplitude of the weights produced by

[8]Note that LMS is not necessarily the best estimator. In fact, the cognitive radio could maintain a bank of multiple estimators for each channel, and then select the best one.

[9]The cognitive transmitter will transmit an independent vector at each iteration. In this iterative procedure, it is not possible for the cognitive radio to mimic a temporal convolution in space by shifting the weights across the antennas.

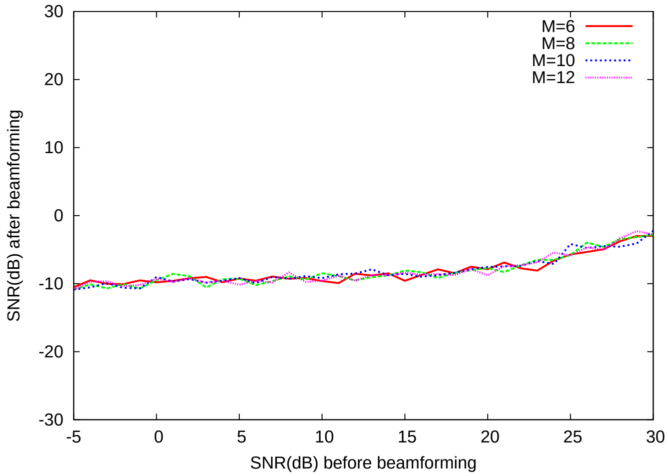

(a)

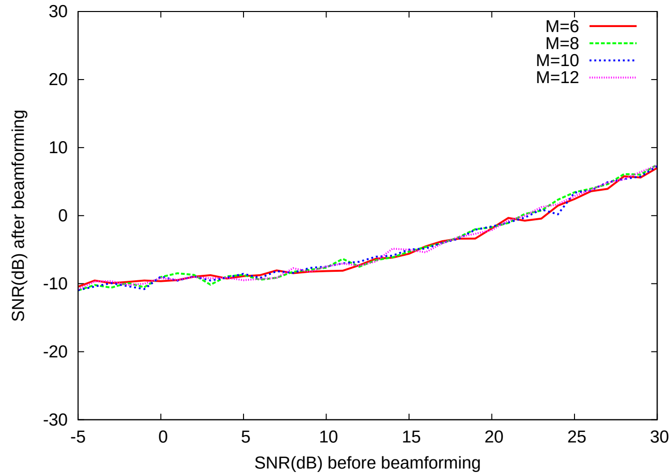

(b)

Fig. 8. (a) The SNR at a primary receiver after beamforming as a function of the SNR before beamforming, for a time-invariant channel and a total of $K = 4$ primary receivers. The curves represent different numbers of antennas at CR_t. (b) An identical experiment with quantization noise added to the beamforming weights.

the optimal beamforming algorithm, to model finite resolution effects in practical implementations. A uniform random variable between ± 6 degrees was added to the phase of each antenna weight, and the amplitude of each weight was scaled by $(1 + \epsilon)$, where ϵ is a uniform random variable in the range $[-0.02, 0.02]$.

We can see that the interference rejection does not depend on the number of antennas M. This is to be expected, since $M > K + 1$ in all cases [16]. We also observe that at low input SNR, the output SNR in the two plots is nearly identical. This means that at low SNR, the dominant error source is the inability of the adaptive filter to accurately estimate the channel. However, for high input SNR, quantization of the beamforming weights increases the output SNR by almost 10 dB. Thus, at high SNR, quantization noise has a significant impact on the performance of the algorithm.

We can also see that both curves are relatively flat. As we vary the input SNR over a range of 35 dB, the output SNR varies by 8 dB (without quantization) or 15 dB (with quantization). We can interpret these results as follows. If the initial SNR is relatively high, then we expect that the estimate of the channel closely tracks the actual channel and the optimal beamformer can achieve very high levels of interference rejection, which is required for a signal with high SNR. On the other hand, when the initial SNR is low, we expect the quality of the LMS estimate to be degraded, which in turn will limit the achievable interference rejection. However, since the SNR is low to begin with, the amount of rejection required to bring the interference power below the noise floor is much smaller!

We also simulated the performance of the iterative, Gram-Schmidt estimation scheme in Section IV-C. We again considered a time-invariant-channel, with $K = 4$ primary receivers and $M = 12$ antennas at the cognitive radio. The results are shown in Figure 9(a), where there is no quantization, and in Figure 9(b), where the same phase and amplitude errors as in the previous experiment were added. It can be seen that the iterative channel estimation procedure does not affect the quality of the interference rejection.

The accuracy of the channel estimate could be further improved by extending the estimation period and averaging the noise. This technique, however, is only effective for time-invariant channels and cannot be used with wireless channels, which are usually time varying because the environment is constantly changing. The rate of change (or the coherence time) of the channel places a limit on the estimation accuracy. To quantify the impact of channel variation, we repeated the first simulation with the channel taps varying in time. The results are shown in Figure 10. The horizontal axis denotes the iteration number, while the vertical axis shows the SNR after beamforming. We used a fixed initial SNR of 30dB, with $M = 10$ antennas and $K = 4$ interferers. A uniform random variable between ± 3 degrees was added to the phase of each antenna weight and the amplitude of each weight was scaled by $(1 + \epsilon)$, with ϵ a uniform random variable in the range $[-0.02, 0.02]$, to model

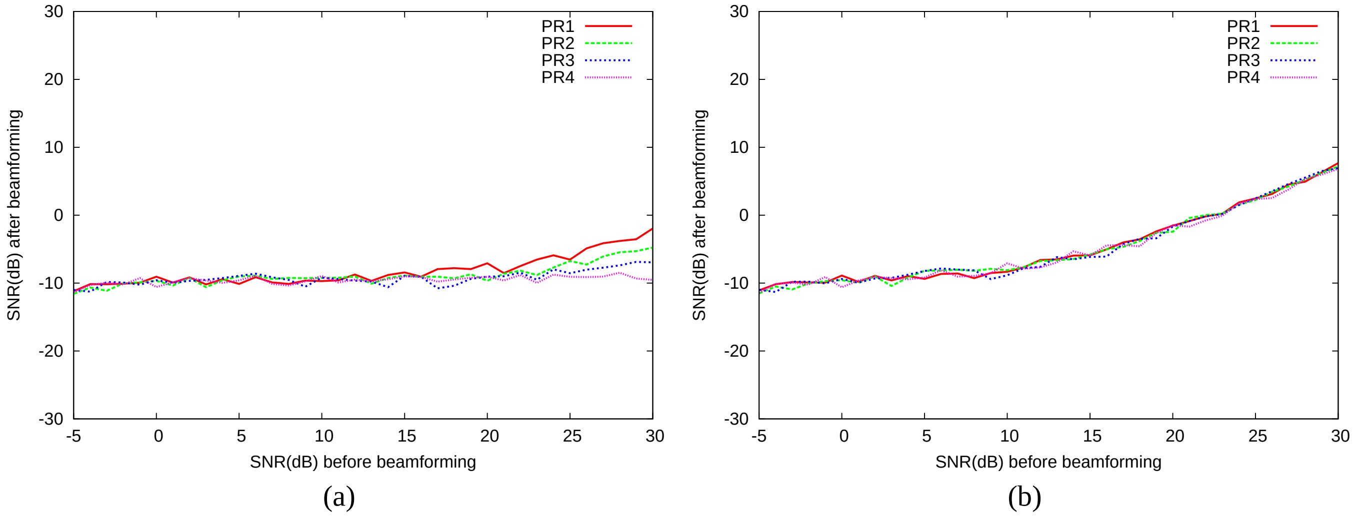

Fig. 9. (a) The SNR at 4 different primary receivers, when the Gram-Schmidt method is used to iteratively estimate the channels, as a function of the input SNR. The cognitive radio has $M = 12$ antennas. (b) An identical experiment with quantization noise added to the beamforming weights.

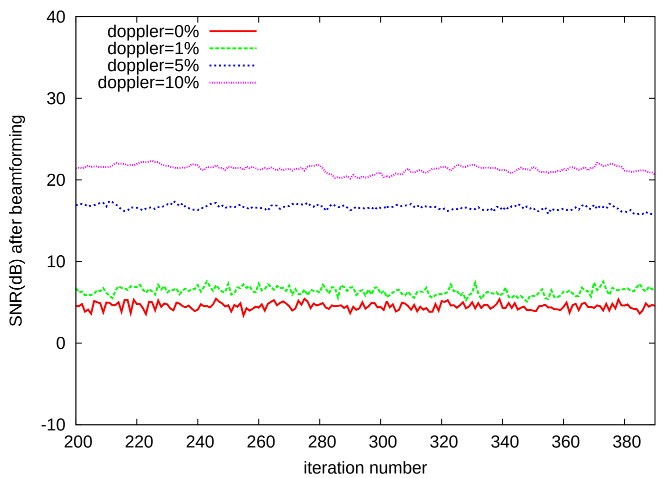

Fig. 10. The SNR after beamforming as a function of time, for a time-varying channel. The SNR before beamforming is 30 dB. The curves represent different doppler rates as a percentage of the LMS update rate.

finite resolution effects. The different curves represent different doppler rates (the inverse of the coherence time), as a percentage of the LMS update rate. The results show that when the doppler rate is 1% of the LMS update, the SNR is only a few dB worse than in the case of a static channel. However, when the doppler rate is 5% or more, it is not possible to estimate the channel sufficiently well and the beam nulling performance is significantly degraded.

V. Conclusion

This paper proposes a new framework for spectrum reuse that relies on collaboration between primary and secondary systems. Secondary radios, assisted by feedback from the primary system, use beam nulling to eliminate interference with the primary receivers. As a demonstrative example, we show how this paradigm can be used to allow secondary systems to reuse the uplink of an OFDMA cellular system. We also demonstrated simple and efficient beamforming (nulling) algorithms that can be used to implement the system. As this application shows, spectrum sharing can be accomplished with little hardware complexity in the secondary radios (as few as two antennas per radio) and little change to the primary infrastructure. Both factors are necessary to the success and future adoption of a system.

Furthermore, in addition to the low deployment and transition costs, this new paradigm presents advantages to both primary and secondary users. The primary system is in full control of how and when the spectrum is shared, and thus can provide guarantees on the service degradation experienced by its users. Cooperation also gives the secondary the benefit of more effective and predictable spectrum access than in opportunistic paradigms. Finally, in this framework spectrum sharing is driven by economic forces and does not require government regulation.

There are still a number of open questions which require further investigation. For example, the impact of various properties of primary systems (e.g. size, scale, topology, protocols, etc.) on the design of secondary systems and the cooperation framework must be analyzed. Similarly, practical beamforming and nulling algorithms for more complex applications must be developed. Our goal in this paper is to introduce the cooperation framework as a new direction in cognitive radio research and present guidelines for designing future systems.

References

[1] R. W. Broderson, A. Wolisz, D. Cabric, S. M. Mishra, and D. Willkomm, "Corvus: A cognitive radio approach for usage of virtual unlicensed spectrum," Tech. Rep., 2004.

[2] T. Erpek, M. Lofquist, and K. Patton, "Spectrum occupancy measurements: Loring Commerce Centre, Limestone, Maine, September 18-20, 2007," Shared Spectrum Company, Tech. Rep., 2007. [Online]. Available: http://www.sharedspectrum.com/measurements/download/Loring_Spectrum_Occupancy_Measurements_v2_3.pdf

[3] A. Sahai, R. Tandra, S. M. Mishra, and N. Hoven, "Fundamental design tradeoffs in cognitive radio systems," in *Proc. of Technology and Policy for Accessing Spectrum*, August 2006.

[4] FCC, "Spectrum policy task force report," ET Docket No. 02-155, November 2002.

[5] The New America Foundation and The Shared Spectrum Company, "Dupont circle spectrum utilization during peak hours," Tech. Rep., June 2003. [Online]. Available: http://www.newamerica.net/files/archive/Doc_File_183_1.pdf

[6] FCC, "FCC 00-402," November 2000. [Online]. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-00-402A1.pdf

[7] A. Ghasemi and E. S. Sousa, "Collaborative spectrum sensing for opportunistic access in fading environments," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, November 2005.

[8] S. M. Mishra, R. Tandra, and A. Sahai, "Coexistence with primary users of different scales," in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, April 2007.

[9] R. Tandra and A. Sahai, "SNR walls for signal detection," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 1, pp. 4–17, February 2008.

[10] S. M. Mishra, A. Sahai, and R. Brodersen, "Cooperative sensing among cognitive radios," in *IEEE International Conference on Communications (ICC)*, June 2006.

[11] R. Mudumbai, J. Hespanha, and U. Madhow, "Scalable feedback control for distributed beamforming in sensor networks," in *IEEE International Symposium on Information Theory (ISIT)*, September 2005.

[12] R. Mudumbai, B. Wild, U. Madhow, and K. Ramchandran, "Distributed beamforming using 1 bit feedback: from concept to realization," in *Allerton Conference on Communication, Control, and Computing*, September 2006.

[13] D. Tse and P. Viswanath, *Fundamentals of Wireless Communications*. Cambridge University Press, 2005.

[14] D. Manolakis, V. Ingle, and S. Kogon, *Statistical and Adaptive Signal Processing: Spectral Estimation, Signal Modeling, Adaptive Filtering and Array Processing*. McGraw-Hill, 2000.

[15] M. Hayes, *Statistical Digital Signal Processing and Modeling*. Wiley, 1996.

[16] H. V. Trees, *Optimum Array Processing*. Wiley-Interscience, 2002.