This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the WCNC 2009 proceedings.

Multi-Antenna Interference Cancellation Techniques for Cognitive Radio Applications

Omar Bakr[†] Mark Johnson Raghuraman Mudumbai Kannan Ramchandran[*]

***Abstract*— This paper presents a practical method for using multi-antenna radios to cancel interference in cognitive radio systems. Under this method, secondary radio transmitters use beamforming techniques to find antenna weights that place nulls at the primary receivers, and secondary radio receivers use adaptive techniques to decode in the presence of interference from primary users. As an example, we show how this scheme can be leveraged to effectively reuse the uplink band of a cellular network. However, estimating the channel responses, without causing interference and without requiring significant modifications to legacy systems, is a challenging problem. We provide an iterative method for accurate channel estimation in frequency division duplexed networks, where the uplink is independent of the downlink.**

I. Introduction

Spectrum is becoming an increasingly scarce and valuable resource. While all bands below 3 GHz have been allocated [1], the demand for new wireless applications continues to grow exponentially. To enable the deployment of new systems, as well as to increase the number of users served by existing systems, a variety of spectrum sharing and reuse schemes have been proposed. One such technology is cognitive radio [2], which is based on radios that can operate in multiple frequency bands and dynamically adapt their transmissions to their environment. These cognitive radios (the secondary system) must be able to structure their transmissions so that they cause negligible interference to the legacy users (the primary system). It is desirable for the secondary system to be as transparent as possible to the primary.

Much prior research on cognitive radio technology has focused on secondary systems that sense the presence or absence of primary radios, and opportunistically reuse the spectrum [1], [3]. The performance of these systems can be improved by collaborative sensing among the secondary radios [4], [5]. However, fundamental limitations on the detection of signals in low SNR environments pose a significant obstacle to the opportunistic reuse paradigm [6].

In this work, we explore practical algorithms for spectrum reuse that rely on collaboration between the primary and secondary systems, and multi-antenna beamforming [7]. Multiple antenna radios are a natural fit for the cognitive radio application, because they offer the possibility of choosing antenna weights that eliminate interference with the primary network. Ideally, we want to select weights so that the secondary transmitter has "nulls" in the directions of the primary receivers, and the secondary receiver is able to cancel interference from primary transmitters. We show in this paper that simple adaptive interference cancellation techniques can be extended to determine the weights of the secondary receiver. However, determining the transmitter weights is more challenging. The secondary transmitter first needs to estimate the unknown channels to the primary receivers, without generating interference in the learning phase. We propose an iterative approach where the secondary transmitter adjusts its weights to null out one primary receiver at a time.

The information-theoretic limits of multi-antenna spectrum reuse systems have been studied in [8], [9], and a game-theoretic approach for power control in multiple antenna cognitive radios was given in [10]. To the best of our knowledge, this paper is the first attempt to explore beam-nulling using multiple antennas to perform interference cancellation for cognitive radio systems.

The rest of the paper is organized as follows. We begin by giving an overview of the collaborative reuse framework in Section II. In Section III, we show how this paradigm can be applied to reusing the uplink channel of a frequency division duplexed cellular network, and describe the problem of estimating channel responses without causing interference to primary users. We present an iterative solution to this problem in Section IV. Numerical simulations of our algorithms are provided in Section V and a discussion of future work is given in Section VI.

II. Cooperative Spectrum Reuse Framework

We have recently proposed a general framework for cooperative spectrum reuse based on beamforming and beam nulling [7]. As an example, we demonstrated how the uplink channel of a cellular network could be reused via this scheme. A multi-antenna cognitive radio (CR) node registers with the primary network as though it were a regular mobile client (a so called *dual citizen*). The cognitive radio then transmits a pilot signal, and nearby base stations feed back their received signals. Based on this feedback, the cognitive radio can estimate the channels and choose a set of antenna weights that produce a radiation pattern with nulls in the directions of the base stations. The cognitive radio then uses those weights to reuse

O. Bakr, M. Johnson, and K. Ramchandran are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA, 94720 (email: {ombakr, mjohnson, kannanr}@eecs.berkeley.edu).

R. Mudumbai is with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, 93106 (email: raghu@ece.ucsb.edu).

[†]Omar Bakr's research is sponsored by a fellowship from King Abdullah University of Science and Technology.

[*]Kannan Ramchandran's research is supported by the National Science Foundation under grants CCF-0729237 and CCF-0635114.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the WCNC 2009 proceedings.

the entire uplink channel, without causing any interference at the primary receivers (base stations).

This framework provides two significant advantages. First, the primary system is in full control of the spectrum sharing, and retains the ability to silence the secondary radios if it experiences unacceptable interference levels. Also, the secondary system receives more predictable spectrum access than under opportunistic reuse paradigms.

Cellular uplink bands can be effectively reused under this framework because the receivers are sparsely distributed and have static locations. Thus, the CR nodes need only null a relatively small number of directions, and the temporal channel variations are smaller than in the downlink channel, where the receivers (the mobiles) are moving. In addition, it is much more practical to modify the base stations to provide feedback than to modify all of the mobile units.

In this work, we develop an algorithm for efficiently estimating the channels and computing the antenna weights, without causing interference to the primary system. This method is designed for use in frequency division duplexing (FDD) systems, which are predominant in many countries due to legacy compatibility reasons. The performance of the secondary system can be further improved if the CR receiver cancels interference from the primary transmitters (the mobiles), using knowledge of the CR transmitter's training sequence.

III. Spectrum Reuse in FDD Networks

Fig. 1 illustrates the basic channel reuse scenario. The cognitive transmitter CR_t wishes to transmit to cognitive receiver CR_r, while placing nulls in the directions of the primary receivers $PR_1, \ldots PR_K$. The CR transmitter has $M > K$ antennas, and the primary radios each have a single antenna. As we are considering an FDD system, the primary radios receive on this band, but transmit on a different band.

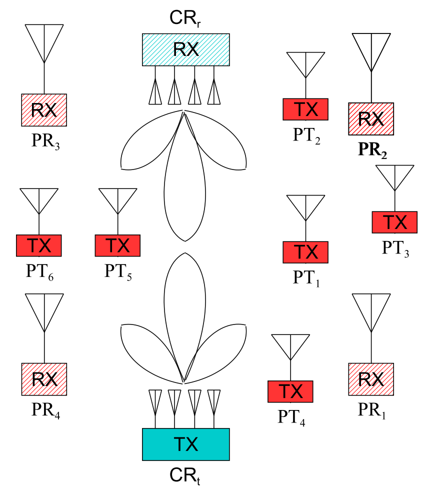

Fig. 1. A cognitive transmitter must simultaneously estimate the channels to four primary receivers in order to reuse the uplink channel. The cognitive radio receiver needs to cancel interference from the primary transmitters in order to receive data from the cognitive radio transmitter.

We assume that both systems are narrowband, so that the channel responses are represented by a single tap in the time domain[1]. The complex channel responses from CR_t to CR_r is denoted by $\mathbf{h}_r = [h_{r,1} \ldots h_{r,M}]$ and the channel from CR_t to PR_i, $1 \leq i \leq K$, is denoted by $\mathbf{h}_i = [h_{i,1} \ldots h_{i,M}]$.

When CR_t transmits a symbol $x[n]$, it is premultiplied by a complex weight c_j^* at antenna j. The received signal at 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 antenna weight vector and $\nu[n]$ is white noise. We assume for simplicity that the primary and secondary systems use the same symbol duration.

If the channel responses are known, the optimal weight vector can be found by projection, as discussed in Section IV-D. In practice, however, the transmitter must estimate the channels. In an FDD system, channel reciprocity may not hold. Thus, learning the uplink channel requires the cognitive radio to transmit on the uplink and receive feedback from the primary receivers. Because CR_t is a dual-citizen, it has been allocated a small slice of the uplink that can be used for this purpose[2]. Note that since the cognitive radio is also a member of the primary system, it does not have to limit its power when using the allocated channel resource.

Since different base stations may allocate different channel resources to the cognitive transmitter, the process of channel estimation can lead to unacceptable interference. Consider a simple example with two primary receivers that have allocated distinct frequencies f_1 and f_2 to the secondary radio. Without loss of generality, assume that the channel to PR_1 is stronger. In order to communicate with PR_2, the cognitive radio must transmit a signal at frequency f_2 at a power large enough to reach both receivers. However, PR_1 has most likely assigned frequency f_2 to another client. Thus, while attempting to estimate $\mathbf{h}_2$, CR_t causes unacceptable interference at PR_1.

This is the fundamental challenge in reusing the uplink channel of an FDD cellular network. When the various primary receivers have allocated distinct channel resources to the secondary radio, it must be able to estimate the unknown channels without causing interference in the primary system. In Section IV, we present an algorithm for estimating the channels that does not generate interference, regardless of whether the primary receivers have assigned different channel slices to CR_t.

IV. Channel Estimation and Interference Cancellation

The key idea behind our interference cancellation technique is for the cognitive transmitter to obtain channel estimates through feedback from the primary receivers and use these estimates to compute the appropriate antenna weights. The

[1] In wideband systems, OFDM can be used to convert the channel into parallel, narrowband subchannels.

[2] The secondary radio will be allocated a time slot, subcarrier, hopping sequence, etc., depending on the physical layer of the primary network.

central difficulty, as stated above, lies in the requirement that the secondary must not generate interference in the process of channel estimation.

A. Naive Channel Estimation

The intuitively simplest estimation scheme would have the cognitive radio transmit a pilot signal from each antenna separately, and then use feedback to compute the channel responses separately. This is equivalent to using a weight vector $\mathbf{c}$ with $M-1$ components equal to zero. Such a framework would emulate receive beamforming, where a multi-antenna receiver can learn the elements of the channel vector individually.

However, the naive scheme suffers from two shortcomings. Principally, transmitting from a single antenna results in omnidirectional radiation, and thus CR_t interferes with all primary receivers when estimating a single channel $\mathbf{h}_i$. Also, by only using one antenna at a time, the SNR at the receiver may not be high enough to enable accurate channel estimation.

B. System Identification via Adaptive Filtering

Adaptive filtering can also be used to estimate the channels. However, like the naive method, such a scheme may generate unacceptable interference, as detailed here. Fig. 2 shows how the channels from a multi-antenna cognitive radio to a single primary receiver can be estimated by posing the problem in the system identification framework. At time n, the cognitive radio uses a random vector of antenna weights $\mathbf{w}[n]$[3]. These antenna weights are also used as the input to an adaptive filter with taps $\mathbf{g}[n]$. The receiver feeds back the received signal $y[n]$, and the adaptive filter coefficients are updated using a standard algorithm such as LMS. When the step size in the filter is chosen appropriately, the filter coefficients $\mathbf{g}[n]$ will approximate the channel $\mathbf{h}$. The cognitive transmitter maintains one filter for each channel that it is estimating.

This architecture requires very little modification to the primary network. Observe that the primary receiver merely feeds back the signal that it detects, and does not perform any computation.

C. Iterative Estimation of Multiple Channels

When the various primary receivers have allocated different channel resources to CR_t, the basic system identification scheme in Section IV-B leads to unacceptable interference. With multiple receivers, let $\mathbf{w}_i[n]$ denote the weight vector used on the channel resource assigned by PR_i. The key observation is that by independently choosing the elements of $\mathbf{w}_i[n]$, the secondary radiates omnidirectionally. Thus, while transmitting on the channel slice assigned by PR_1, the secondary radio interferes with the other primary receivers.

In order to control this interference, we modify the original algorithm. Again consider the scenario in Fig. 1, and assume without loss of generality that the strongest uplink channel from CR_t is to PR_1, the second strongest channel is to PR_2, etc. The cognitive radio initially transmits at power low

[3]The components of $\mathbf{w}[n]$ are independent in time and space, and can be sampled from any distribution (e.g. Gaussian, Bernoulli, etc.).

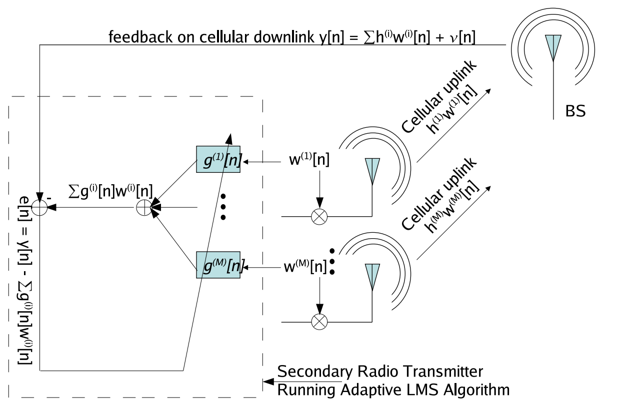

Fig. 2. The transmitter sends a random pilot sequence of weights, and employs an adaptive filter to estimate the unknown channel gains.

enough that it can only reliably communicate with the first receiver[4]. The algorithm described in Section IV-B is then used to compute $\hat{\mathbf{h}}_1$, an estimate of $\mathbf{h}_1$. This will cause very little interference at the other primary receivers, since the channels from CR_t to each of them are assumed to be weaker than $\mathbf{h}_1$.

The cognitive radio then increases its transmit power, so that it communicates only with PR_1 and PR_2. At time n, it first generates a white vector $\mathbf{w}_2[n]$, and then finds the component orthogonal to $\hat{\mathbf{h}}_1$:

$$\tilde{\mathbf{w}}_2[n] = \mathbf{w}_2[n] - \langle \mathbf{w}_2[n], \hat{\mathbf{h}}_1 \rangle \cdot \hat{\mathbf{h}}_1$$

Thus, when CR_t uses $\tilde{\mathbf{w}}_2[n]$ to transmit on the channel resource assigned by PR_2, little or no power is received at PR_1, as long as $\hat{\mathbf{h}}_1$ is an accurate estimate of $\mathbf{h}_1$. Let us express the channel to PR_2 as $\mathbf{h}_2 = \mathbf{h}_2^{\|} + \mathbf{h}_2^{\perp}$, where $\mathbf{h}_2^{\|}$ is the component parallel to $\mathbf{h}_1$ and $\mathbf{h}_2^{\perp}$ is the component orthogonal to $\mathbf{h}_1$. By choosing the training sequence in this manner, the adaptive filter coefficients $\mathbf{g}_2[n]$ will converge to $\hat{\mathbf{h}}_2$, an estimate of $\mathbf{h}_2^{\perp}$.

This algorithm then proceeds in an iterative manner. When estimating $\mathbf{h}_3$, the cognitive radio generates a white random vector $\mathbf{w}_3[n]$, and then finds the component orthogonal to both $\hat{\mathbf{h}}_1$ and $\hat{\mathbf{h}}_2$. Using these weights to send a pilot on the channel resource assigned by PR_3 will not cause interference at either PR_1 or PR_2. Our scheme, which iteratively learns the orthogonal component of each channel, can be viewed as computing the Gram-Schmidt basis of the set of channel vectors.

In practice, there will be frequency offsets between the cognitive radio and the primary receivers, which will result in the baseband channel responses $\mathbf{h}_i$ varying in time. However, because CR_t is estimating the channels in parallel, it can estimate each frequency offset individually and compensate for them in the adaptive filters.

[4]This could be accomplished by having CR_t start with a very low transmit power, which is gradually increased until one base station responds.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the WCNC 2009 proceedings.

With careful choice of beamforming weights, the cognitive radio transmitter can simultaneously use the channels (e.g., hopping sequences) allocated by the primary receivers (base stations) for both channel estimation and transmitting data to the cognitive radio receiver. As long as there is sufficient spatial separation between the weight vectors used for data transmission and those used for channel estimation, the cognitive radio receiver can reliably decode the bits from the cognitive radio transmitter. This allows the cognitive radios to communicate using the entire uplink (including the channels allocated by the primary system), while continuously adapting to changes in the environment.

D. Computation of Antenna Weights

Once the channels have been estimated, the optimal antenna weights $\mathbf{c}$ can be found by projecting $\mathbf{h}_r$ onto the subspace that is orthogonal to the space spanned by the set $\{\mathbf{h}_i, 1 \leq i \leq K\}$. Without loss of generality, we will restrict $\mathbf{c}$ to have unit magnitude. The resulting weights perfectly null the directions of all primary receivers. Furthermore, when the dimension of the signal $\mathbf{h}_r$ (the number of antennas M) is larger than the dimension of the interference subspace (K), projection will result in little sacrifice of desired signal strength, with high probability.

We define the interference rejection of this beamformer toward primary receiver 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 weights $c_j = 1/\sqrt{M},\ 1 \leq j \leq M$). Observing that the magnitude of the channel response with beamforming is $|\mathbf{c}^h \mathbf{h}_i|$, and without beamforming is $\left|\frac{1}{\sqrt{M}} \sum_{j=1}^{M} h_{i,j}\right|$, we see that the interference rejection is equal to

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

E. Interference cancellation at the CR receiver using adaptive filtering

Up to this point we have addressed the problem of determining the antenna weights for the cognitive transmitter CR_t. We now describe a simple adaptive algorithm to allow the cognitive receiver CR_r to minimize interference from the primary system. Let $d[n]$ be the symbols sent by CR_t to CR_r[5], and $d_i[n]$ the unknown symbols from the primary transmitter i. We assume $E[|d[n]|^2] = E[|d_i[n]|^2] = 1$. The input signal at CR_r's receive array is:

$$\mathbf{y}[n] = \mathbf{h}_r d[n] + \sum_{i=1}^{K} \mathbf{h}_i d_i[n] + \mathbf{v}[n] \tag{1}$$

where $\mathbf{v}[n]$ is additive white noise and $\mathbf{h}_i$ represents the combined gain of the channel and the transmit beamforming weights.

[5] $d[n]$ could be either a known training sequence or obtained in a decision directed manner (based on receiver decisions).

If the receiver uses the unit-norm weight vector $\mathbf{c}$, and σ_v^2 is the variance of each component of $\mathbf{v}[n]$, then the resulting signal to interference and noise ratio (SINR) is given by

$$\text{SINR} = \frac{|\mathbf{c}^h \mathbf{h}_r|^2}{\sum_i |\mathbf{c}^h \mathbf{h}_i|^2 + \sigma_v^2} \tag{2}$$

We seek the MMSE weight vector $\mathbf{c}_{MMSE}$ that maximizes the SINR:

$$\mathbf{c}_{MMSE} \triangleq \arg\min_{\mathbf{c}} |e[n]|^2, \text{ where } e[n] \triangleq \mathbf{c}^h \mathbf{y}[n] - d[n] \tag{3}$$

Earlier work on CDMA systems has shown that adaptive algorithms such as the LMS, normalized-LMS (NLMS) and RLS algorithms are able to achieve interference cancellation and also are immune to the near-far problem. The convergence properties of the adaptive algorithms have been studied in great detail in the literature [11]. We only note that there is always a tradeoff between the ability of the algorithms to track time-variations and the mean squared error achievable, and this tradeoff can be controlled by adjusting algorithm parameters (e.g., the step-size in the LMS and the forgetting factor in the RLS algorithms). Intuitively a slower convergence allows less noisy updates, but fails to track fast channel variations. As seen from (1), the carrier frequency offsets between the different transmitters appears as a time-varying effect on the channels; since these offsets can be quite substantial[6], this imposes a limit on the performance of these algorithms.

One approach to dealing with time-variations is to use a differential variation of the MMSE criterion in (3). This DMMSE criterion was introduced in [12] and is defined as:

$$\mathbf{c}_{DMMSE} \triangleq \arg\min_{\mathbf{c}} \left|\mathbf{c}^h \mathbf{y}[n-1] d[n] - \mathbf{c}^h \mathbf{y}[n] d[n-1]\right|^2,$$
$$\text{subject to } E[|\mathbf{c}^h \mathbf{y}[n]|^2] = 1 \tag{4}$$

The intuition behind this criterion is that even the most rapidly varying channels can be assumed to be constant for two consecutive symbols. Therefore, by minimizing the error in tracking the *ratio* of two consecutive symbols, the algorithm is spared the task of explicitly tracking the channel variations. It was shown in [12] that the DMMSE criterion in (4) converges to the MMSE weight $\mathbf{c}_{MMSE}$ under mild assumptions. Furthermore, the standard adaptive RLS algorithm can be easily adapted for (4). Our simulation results show that the adaptive DMMSE algorithm achieves almost 10 dB higher SINR compared to the standard RLS or LMS algorithms when the frequency offsets and Doppler spreads are significant.

V. Simulation results

We simulated the performance of the scheme in Section IV-C in order to evaluate it in realistic environments. We first considered the case of a time-invariant channel. The plot in Fig. 3(a) shows the SNR at one primary receiver after

[6] Since the cognitive radio receiver only receives the aggregate signal from different primary transmitters as well as the cognitive radio transmitter, it is difficult to compensate for the different frequency offsets even if they are known a priori.

beamforming as a function of the SNR before beamforming[7]. There were a total of $K = 4$ primary receivers in the system, each estimated at the same input SNR, and the cognitive transmitter has $M = 12$ antennas. Note that the SNR before beamforming is the SNR at which the adaptive filter operates while estimating the channel. When $M > K$, increasing the number of antennas does not reduce the SNR at the primary receivers [13], as long as the total transmit power is constant (independent of M)[8]. Fig. 3(b) shows the output SNR when quantization noise is added to the phase and amplitude of the weights produced by 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 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 the 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 18 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 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 channel 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.

The accuracy of the channel estimates 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, since 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 this simulation with the channel taps varying in time. The results are shown in Fig. 3(c). The horizontal axis denotes the iteration number, while the vertical axis shows the SNR after beamforming. We used a fixed input 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 account for errors caused by noise and limited precision. The different curves represent different Doppler spreads (the inverse of the coherence time), as a percentage of the LMS update rate. The results show that when the Doppler spread is 1% of the LMS update, the SNR is only a few dB worse than in the case of a time-invariant channel. However, when the Doppler spread is 5% or more, it is not possible to track the channel sufficiently well, and the beam nulling performance is significantly degraded.

We have also simulated the performance of the normalized LMS (NLMS) and DMMSE algorithms for receive beamforming, described in Section IV-E. In both cases, we simulated a cognitive radio receiver with 8 antennas and 5 interferers. The desired pilot signal from the cognitive transmitter is 20 dB below each interferer (thus 27 dB below the interference power, given that there are 5 interferers) and 10 dB above the noise floor. Therefore, the maximum possible output SINR, achieved when all interferers are perfectly canceled, is about 19dB. Fig. 4(a) shows the output SINR as a function of time, when the secondary receiver uses the DMMSE algorithm and the channel is time-invariant. Results of an identical experiment, using the NLMS algorithm at the receiver, are shown in Fig. 4(b). We see that DMMSE converges faster than NLMS, and achieves a better output SNR (closer to the optimum).

Fig. 4(c) shows the performance of the two algorithms with a time varying channel. For the DMMSE algorithm, the Doppler spread is 0.2% of the symbol rate and the frequency offset is 20%[9]. Note the loss in performance caused by the channel variation. For the NLMS algorithm, the Doppler spread is 0.03% and there is no frequency offset. Observe that NLMS is much more sensitive to channel variation. A more comprehensive analysis of how factors such as quantization, Doppler spreads, and frequency offsets affect the performance of the adaptive algorithms will be the subject of future work.

VI. Conclusion

In this work, we have presented an efficient array processing algorithm that enables interference cancellation in cognitive radio systems using the framework discussed in [7]. In order to find a set of antenna weights with nulls in the directions of primary receivers, a secondary radio must learn the relevant channel responses. The algorithm presented here accomplishes this task without generating interference in the primary and without necessitating coordination among the primary receivers. One of the major advantages of this scheme is that only minor modifications to the primary system are required.

[7]Since the secondary system wants to prevent interference to the primary, the goal of the beamformer is to minimize this output SNR.

[8]Note, however, that increasing the number of antennas does increase the gain to the secondary receiver. Thus, the transmitter can potentially reduce the total transmit power, which would reduce interference to the primary receivers, without affecting the performance of the secondary system.

[9]The frequency accuracy of most crystal oscillators found in commodity radios is of the order of tens of parts per million (ppm). This results in an average frequency offset of the order of tens of *KHz* when the carrier frequency is of the order of *1 GHz*. If the symbol rate is *1 MHz*, then the frequency offset will be several percent as a fraction of the symbol rate.

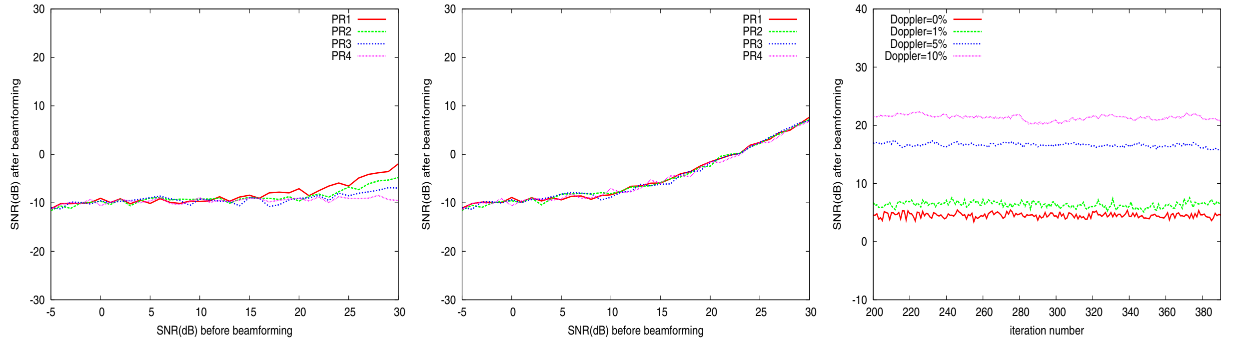

(a) (b) (c)

Fig. 3. (a) The SNR at 4 different primary receivers, when the Gram-Schmidt method is used to iteratively estimate a time-invariant channel, as a function of the input SNR. (b) An identical experiment with quantization noise added to the beamforming weights. (c) The SNR after transmit beamforming as a function of time, for a time-varying channel. The SNR before transmit beamforming is 30 dB. The curves represent different Doppler spreads as a percentage of the LMS update rate.

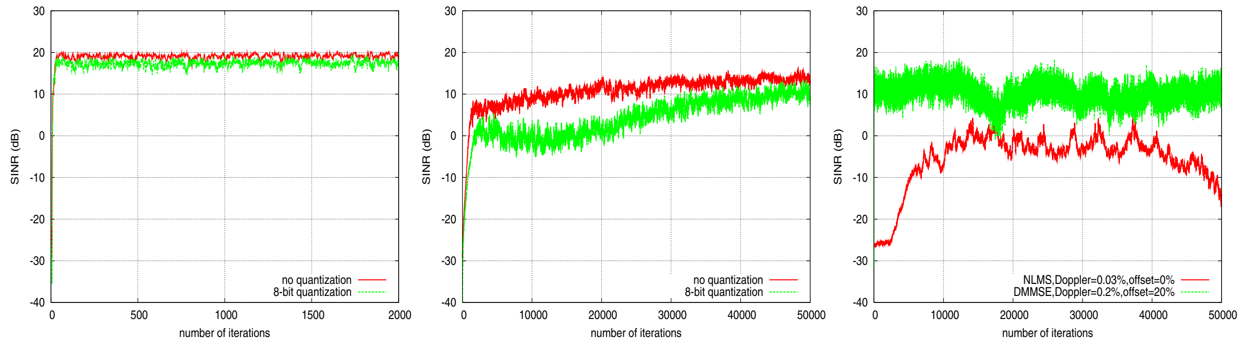

(a) (b) (c)

Fig. 4. (a) The SINR after transmit beamforming as a function of time for a time-invariant channel, with the DMMSE algorithm used at the receiver. (b) An identical experiment, using the NLMS algorithm at the receiver. (c) The SINR after beamforming as a function of time for time varying channels. For the DMMSE algorithm, the Doppler spread is 0.1% of the symbol rate and the frequency offset is 10%. For the NLMS algorithm, the Doppler spread is 0.03% and there is no frequency offset.

Future work in this area includes deriving algorithms that can leverage multiple antennas at the primary receivers and the cognitive receiver, as well as investigating practical adaptive filtering algorithms that boost the interference rejection.

References

[1] 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.

[2] J. Mitola and G. Q. Maguire, “Cognitive radio: making software radios more personal,” *IEEE Personal Communications [see also IEEE Wireless Communications]*, vol. 6, no. 4, pp. 13–18, Aug 1999.

[3] 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.

[4] 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.

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

[6] 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.

[7] O. Bakr, M. Johnson, B. Wild, and K. Ramchandran, “A multi-antenna framework for spectrum reuse based on primary-secondary cooperation,” in *IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, October 2008.

[8] R. Zhang and Y.-C. Liang, “Exploiting multi-antennas for opportunistic spectrum sharing in cognitive radio networks,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 1, pp. 88–93, February 2008.

[9] O. O. Koyluoglu and H. E. Gamal, “On the utility of frequency reuse in cognitive radio channels,” *IEEE International Symposium on Information Theory*, pp. 2161–2165, June 2007.

[10] E. Baccarelli, M. Biagi, C. Pelizzoni, and N. Cordeschi, “Multi-antenna cognitive radio for broadband access in 4G-WLANs,” in *ACM International Workshop on Mobility Management and Wireless Access (MobiWac)*, October 2007.

[11] S. S. Haykin, *Adaptive filter theory*, 3rd ed. Prentice Hall, 1996.

[12] U. Madhow, K. Bruvold, and L. J. Zhu, “Differential MMSE: A framework for robust adaptive interference suppression for DS-CDMA over fading channels,” *IEEE Transactions on Communications*, vol. 53, no. 8, pp. 1377–1390, Aug. 2005.

[13] O. Bakr and M. Johnson, “Impact of phase and amplitude errors on array performance,” EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2009-1, January 2009.