Link Rendezvous Protocol for Cognitive Radio Networks

Brent Horine and Damla Turgut
School of Electrical Engineering and Computer Science
University of Central Florida
Orlando, FL 32816-2362
{bhorine,turgut}@eecs.ucf.edu

Abstract—In an opportunistic dynamic spectrum access environment, individual nodes sense the local spectrum and choose their operating frequencies and bandwidth in collaboration with the other participating nodes. To make their spectrum access decisions, the nodes need to communicate with the existing nodes operating in the area.

In this paper we propose an approach which establishes this first connection with a minimum risk of interference. We propose a link rendezvous strategy which relies on frequency domain decision statistics. Nodes wishing to join the network are emitting and scanning for a simple carrier with a small number of sidetones. We describe the strategy in the context of collaborative spectrum sensing. To validate our approach, we describe a series of experiments using the GNU Radio software defined radio toolkit. We show that an attention signal of length equivalent to a single FFT frame can be detected in a high noise environment using two sidetones.

Index Terms—Dynamic spectrum allocation, link rendezvous and evolution

I. INTRODUCTION

Some cognitive radio (CR) applications require the ability to form networks with minimal or no support from existing infrastructure. In particular, military operations often occur in hostile territory where existing infrastructure is not usable for many reasons. Additionally, covert operations require minimal RF signatures and low probability of intercept. These issues argue against establishing a conventional infrastructure. First responders in disaster scenarios face similar challenges. Existing infrastructure is often destroyed in the disaster. In this case, users desire to establish a network as soon as possible and to evolve that network as more communication assets are deployed to the scene. Self configuring networks are highly advantageous in these scenarios. In both scenarios, it is important to minimize bottlenecks created by rapidly escalating communication activity.

These scenarios require networks to adapt to changing conditions, of which it is impossible to plan for all possibilities. The ability of a network to autonomously evaluate the situation requirements and environment and to synthesize solutions is a significant goal of cognitive radio networks. In order to retain the greatest amount of flexibility from which to derive solutions, many radio parameters are considered variable.

Dynamic spectrum access (DSA) has emerged as a valuable tool in cognitive radio networks. Systems can opportunistically use spectrum in unlicensed bands, or under-utilized spectrum in licensed bands as secondary users, provided that they do not cause harmful interference to licensed, primary users.

In opportunistic architectures, individual radio nodes sense the local spectrum and choose operating frequencies and bandwidths in unused portions of the spectrum bands. Because of mobility, multi-path fading, shadowing, and other complex propagation issues [1], a sensing node's perception of the spectral environment may be quite different from the actual environment immediately surrounding it and different from that at the destination node(s). Unintentional interference is therefore a possibility when sensing nodes are in deep fades and cannot detect the primary user. Once the nodes have established a connection, they can exchange information on the spectral environment so that a more collaborative approach can be used to arrive at superior decisions.

The rest of the paper is organized as follows. We survey the existing literature in Section II. Section III presents the proposed link rendezvous protocol in detail. Experimental study and results are provided in Section IV. We conclude in Section V.

II. RELATED WORK

Polson [2] identifies two main approaches toward link rendezvous. The first assumes some infrastructure which transmits a beacon encoded with information about time and frequency. Cognitive radios request a time-frequency slot for a specified network and provide location and power information. The infrastructure server recommends a frequency and schedules a time for the CR to check back in. The server has an omniscient view of all CRs in the area and can globally optimize its decisions.

As previously noted, infrastructure is not always appropriate. Polson also describes an unaided approach. It relies on the calling node to emit probe signals on a selection of available frequencies. Receiving nodes listen on the set of frequencies which they have determined are available. When the original probe waveform is detected, the receiving node transmits its own probe-acknowledgement waveform. The nodes can then establish a conventional link on that frequency. Given enough time and frequencies, it is expected that a link can be established. Once established, the nodes can exchange information which expands their knowledge of the spectral environment. The collaborative exchange has been studied

by a number of authors. Zhao et. al. propose a system in which nodes self-configure into local groups and establish a common control channel to serve that group [3]. Fringe nodes bridge between groups and provide and overall network connectivity. They also propose a modified MAC which leverages this approach. Ghasemi and Sousa investigated collaborative sensing in fading and shadowing channels [4]. They describe a system which exchanges bit information where each node compares the sensed energy to a threshold and sets the bit value accordingly. When a node receives information from other nodes, it performs a logical OR operation on all of the reporting nodes at a given frequency. Any node reporting a signal therefore dominates the decision. Ganesan and Li describe a relay based approach in a TDMA system in which nodes listen on alternate time slices and pass on information about primary transmitters on their time slice [5].

III. Link Rendezvous Protocol

Conventional radio systems have well defined communication parameters, making the initial establishment of a communication link relatively easy. Often, dedicated calling channels are used to initiate the link, after which the parties negotiate a new set of parameters, e.g. channel, on which to further conduct their communications. While convenient, this approach can degrade the efficiency of spectrum utilization. In congested environments, this signalling channel can also be a bottleneck. Finally, these channels represent a single point of failure in the presence of hostile jamming. Fortunately, dedicated channels are not needed in DSA based cognitive radio networks. CR networks can be designed so that nodes rendezvous with each other based upon the sensed spectral environment provided that the application can withstand some level of delay in initial network setup. This process is known as Link Rendezvous [2]. An algorithm to accomplish such is described here. A heterogeneous mesh network is assumed with some nodes providing application services such as a gateway. A service discovery protocol will be running to assist connecting nodes in finding the services they require.

Figure 1 describes the high level view of the link rendezvous process. This process begins by dividing the spectrum into regularly spaced frequency bins appropriate to the link establishment protocol bandwidth. The node senses the spectrum in each of these bins. The RF energy in each bin is compared to a threshold. Vacant bins are flagged as potential link establishment channels.

Once a certain number of vacant channels are identified, the link requester begins its attempt to connect to or establish a network. It transmits an attention signal with an easily identified spectrum on the channel for T seconds. For example, this spectrum could be a double-sideband signal, amplitude modulated with a set of discrete tones at specific relative amplitudes and offset frequencies. Until the node joins a network, it has a limited view of the spectral environment. In order to minimize unintentional interference, the first attention signal should be emitted with low power and reasonably short duration. Nodes on standby continuously monitor the spectrum

Sense Spectrum

Calculate Parameters
- number of frequencies
- which frequencies
- output power

Emit Attention Signal on Chosen Frequencies

Increase Spectrum Sensing Detection Limits

Response? — No

Yes

Attempt Connection

Attempt again or sense again or listen for others

Acknowledged? — No

Yes

Choose Destination Node

Negotiate Parameters

Choose another node
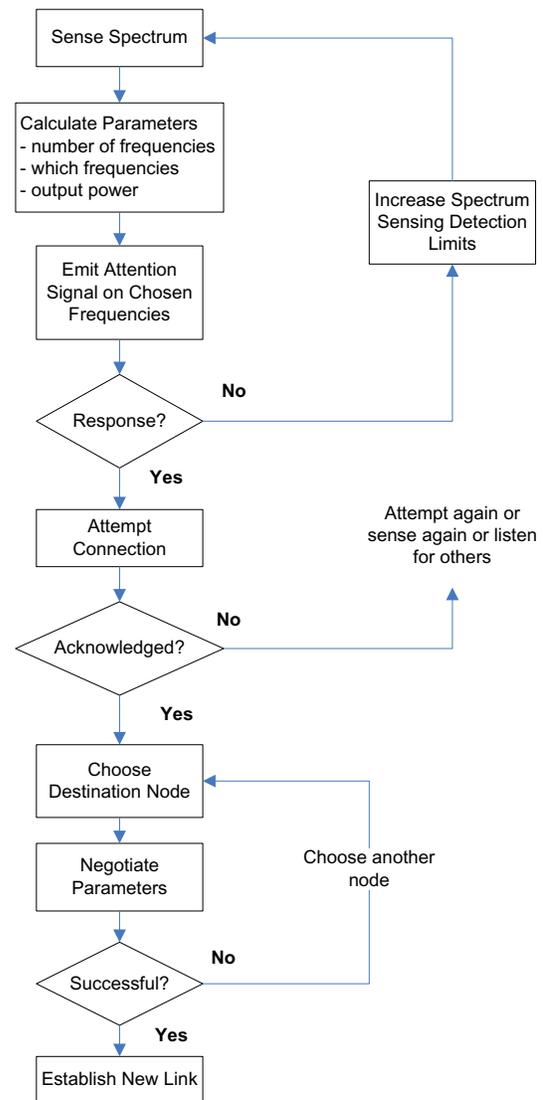
Successful? — No

Yes

Establish New Link

Fig. 1. Initial Connection

for this pattern. SDR technology allows the radio to scan many channels at once. An entire band can be sampled, based upon the performance of the A/D converter and system processing speed. A Fast Fourier Transform (FFT) is performed on the resulting sample stream. A feature detection algorithm then searches for the well defined pattern resulting from the specific modulation parameters. The receiving node collects all of the attention signal occurrences it finds within its scanning range. It then chooses which frequency to use in response. It makes this decision based upon its knowledge about the spectrum occupancy and the signal strength of the incoming signal, along with other channel statistics it may have gathered. It then transmits similar but distinct pattern of its own on the chosen frequency.

After the initial transmission, the calling node switches to a listen mode. It scans all of the frequencies on which it originally transmitted, looking for the reply pattern of sidebands. In dense RF environments, it might receive more than one

reply. The calling node chooses the final frequency on which to connect from the set of responses. This decision may be based upon signal strength or other ranking. The node is finally ready to establish a connection. It broadcasts a connection request code using some reasonably lowest common denominator RF parameters on the chosen frequency.

After transmitting the attention reply signal, idle receivers enter a listen mode for a connection request. Upon detecting a connection request, the receiver transmits a connection response message directly to the originating node. This message is the first unicast message and can include information about the node such as the services it can provide and connection parameter preferences. The originating node chooses to which destination node to connect which finishes the rendezvous process.

The process can fail at several points. No nodes may respond to the attention signal. No nodes may respond to the connection request. The actual connection could fail. The nodes may not have any overlap of perceived vacant bins. If no nodes respond to the attention signal, a corrective action could be to increase the power and duration of the signal. Since this ultimately increases the likelihood that the attention signal will interfere with existing communications, the limits and the rate of increase should be capped by a policy and driven by the application requirements.

The hypothesis that leads to the use of an analog modulated attention signal is that the transmitter can secure the attention of a receiver with a minimal energy, bandwidth and duration. This is achieved primarily by an SDR's ability to monitor multiple channels simultaneously. A digital scheme can also take advantage of this feature of an SDR, but requires acquisition and synchronization times. We now lay the foundation for exploring the tradeoffs associated with detecting this signal in the frequency domain. Note that there are several design questions to answer that also affect the detection limits. Namely, the size of the FFT, the window function applied, and the detection mechanism. Let the transmitted attention signal be

$$s = \cos(\omega_c t) \sum_{j=1}^{J} A_j \cos(\omega_j t) \qquad (1)$$

where j represents the j^{th} modulated tone, ω_c is the carrier frequency, and typically $\omega_j \ll \omega_c$.

An additive white gaussian noise channel is assumed. This results in the received signal

$$r(t) = s(t) + n(t) \qquad (2)$$

$$r(t) = \cos(\omega_c t) \sum_{j=1}^{J} A_j \cos(\omega_j t) + n(t) \qquad (3)$$

where $n(t)$ is the noise.

At this point, either a FFT feature detector or a cyclostationary detector can be used. We choose to perform an FFT and use a correlator in the frequency domain. The FFT step is illustrated in equation (4) where the signal has been sampled and transformed.

$$R(k) = \sum_{n=0}^{N-1} [\cos(\omega_c n) \sum_{j=1}^{J} A_j \cos(\omega_j n) + z(n)]$$
$$\cdot W(n) \exp^{(-j2\pi kn/N)} \qquad (4)$$

where N is the number of samples, k is the frequency parameter, $W(n)$ represents a window function and $z(n)$ is the sampled noise.

Practical implementation issues further constrain the solution. Since the receiver is not synchronized to the start of the attention burst, the frame handed off to the FFT operation might capture as little as one half of the burst. In order to maintain the resolving power, the actual pulse should be twice the minimum length. In turn, this assumes that the sampling and buffering hardware is capable of producing contiguous frames and that the FFT processor is capable of consuming all of the frames in real time.

The decision statistics are ratiometric in nature and can be summarized as in equation (5).

$$H1 : \bigcap_{i=1}^{N} [A_i P_c - tol < S_i < A_n P_c + tol] \qquad (5)$$

where $H1$ represents the hypothesis that an attention signal is present, A_n is the design relative amplitude for the n^{th} sidetone, tol is a tolerance to account for nonlinearities and some noise margin, S_n is the sidetone amplitude, N is the number of sidetones and P_c is the carrier amplitude. The alternate hypothesis, $H0$ is assumed when any one these conditions are not met and indicates that there is not an attention signal present. Wider tol values reduce the false negative rate but increase the false positive rate. False positives decrease spectrum utilization and drain node energy because of the wasted attention response transmissions. They also consume node energy. Narrower tol values decrease the false negative rate, and increase the difficulty and time required to establish a connection. The tolerance value and the number of sidetones are involved in a complex tradeoff with the amount of energy used to transmit the signals and therefore the likelihood of unintentional interference.

IV. EXPERIMENTAL STUDY

A. Testing Environment

In order to test some of these concepts, an experimental study was developed based upon GNU Radio [6]. For this round of experiments, a simulated channel was used. In other words, the signal was not actually transmitted. These tests focused on detecting the attention signal without *a priori* knowledge of the center frequency beyond a band specification and gaining an empirical understanding of how the choice of window function, sidetone amplitude tolerance, and number of sidetones affect the detection limits.

Although GNU Radio is elegantly designed to minimize the amount of custom C++ code required for most applications,

TABLE I
NOMINAL EXPERIMENTAL PARAMETERS

Parameters	Value
Carrier Frequency	1000
Sampling Frequency	2500
Carrier Amplitude	1.0
Offset Frequency 1	12.207
Offset Frequency 2	24.414
Offset Frequency 3	48.828
Relative Sidetone Amplitude	0.20
FFT Length	8192

TABLE II
FALSE POSITIVES VERSUS NOISE BY NUMBER OF SIDETONES

Noise Amplitude	Sidetone 1	Sidetone 2	Sidetone 3
0.01	21	7	7
0.1	48	8	6
0.5	35	2	1
1.0	30	0	0
2.0	33	0	0
3.0	33	0	0

the unique correlation in the frequency domain required some development. In particular, after applying an FFT to the incoming data, the custom module iterates through each sample looking for the correct amplitude response at the sidetone location. The amplitude is measured at a particular bin and is hypothesized to be the carrier amplitude. The amplitude in bins higher and lower by each sidetone's offset is then recorded. The tolerance is calculated relative to the carrier amplitude. The hypothesis that there is an attention signal present at the carrier frequency is true only if each sidetone has an amplitude within the tolerance specification.

The parameters for the test are listed in Table I. The tolerance and gaussian noise amplitude were varied during the individual tests as was the window function applied during the FFT. The frequencies and amplitudes have relative dimensions. The same results are achieved if the units are kHz, MHz, or GHz. The sidetones are chosen so that they are separated by an integer number of bins from the carrier, based upon the sampling frequency and number of FFT points.

Averaging was not utilized. Effectively, the transmission time was equivalent to N/F_s. For convenience, the system built up attention signal from separate RF signals, rather than modulating an RF carrier with the baseband tones. Gaussian noise was added at the source. No filtering was applied, so there is some noise being aliased into the measurement.

B. Results

As described previously, the time-truncation of the attention signal causes a degradation of the frequency resolution and the appearance of leakage. The degradation of the frequency resolution is due to the main lobe widening. If the signal was infinitely long and the analysis was likewise infinite, the carrier and sidetones would appear as infinitely narrow pulses in the frequency domain. Instead, the main lobe and the sidetone lobes are broadened significantly. The broadening is smooth and consistent between the carrier and the sidetones. This results in false positives in frequency bins adjacent to the carrier bin, since the test is based upon a relative amplitude. Fortunately, these false positives appear symmetrically around the true carrier. A post processor could be used to perform a frequency average. The performance as the number of sidetones varies is summarized in Table II. It lists the number of false positives found in a single scan of the spectrum. A

rectangular window was used with 8192 FFT points. As the noise increases, the single sidetone has more false positives. On the other hand, the double and triple sidetone versions manage to properly identify the signal with no false positives. The tolerance was held constant at 0.1. This causes an increase in false positives at the higher signal to noise ratios. An adaptive approach to setting the tolerance would drive the false positive rate to near zero for low levels of noise. Clearly, the double sidetone version performs better than the single sidetone; however, there is little to be gained with the triple sidetone. It should be noted that the single tone performance can be significantly improved at moderate to low noise levels by adjustment of the tolerance value. Additionally, the false positives at the lower noise levels are generally in adjacent bins.

V. CONCLUSIONS

A strategy for link rendezvous has been presented. This approach avoids a dedicated signalling channel, only requiring radios to operate within a common band. The concept minimizes unintentional interference during the rendezvous process by using a very short duration, narrow bandwidth, low power attention signal. The responding nodes begin coordinating the spectrum sensing by responding to the attention signal on a frequency which it interprets as being clear.

The minimum time duration of the attention signal is driven by a number of theoretical and practical implementation parameters. By using ratiometric decision statistics and multiple, independent sidetones, the false positive and false negative error rates can be set, with associated tradeoffs in bandwidth, output power, and time duration. These factors ultimately determine the risk of interference to existing communications.

REFERENCES

[1] D. N. Hatfield and P. J. Weiser, "Property rights in spectrum: Taking the next step," in *DySPAN 2005*, November 2005, pp. 43–55.
[2] J. Polson, "Cognitive radio: The technologies," in *Cognitive Radio Technologies*, B. Fette, Ed. Elsevier, 2006, ch. 4, pp. 153–155.
[3] J. Zhao, H. Zheng, and G. H. Yang, "Distributed coordination in dynamic spectrum allocation networks," in *DySPAN 2005*, November 2005, pp. 259–268.
[4] A. Ghasemi and E. S. Sousa, "Collaborative spectrum sensing for opportunistic access in fading environments," in *DySPAN 2005*, November 2005, pp. 131–136.
[5] G. Ganesan and Y. G. Li, "Cooperative spectrum sensing in cognitive radio networks," in *DySPAN 2005*, November 2005, pp. 137–143.
[6] "Gnu radio." [Online]. Available: http://www.gnu.org/software/gnuradio/