A SPATIAL AND TEMPORAL SPECTRUM SENSING SYSTEM FOR INTERFERENCE AVOIDANCE IN DYNAMIC SPECTRUM ACCESS

Reginald L. Cooper (Carnegie Mellon University, Pittsburgh, PA, USA; rcooper1@andrew.cmu.edu); Kevin C. Borries (Carnegie Mellon University, Pittsburgh, PA, USA; kborries@andrew.cmu.edu); Xiaohui Wang (Carnegie Mellon University, Pittsburgh, PA, USA; xiaohuiw@andrew.cmu.edu); and Daniel D. Stancil (Carnegie Mellon University, Pittsburgh, PA, USA; stancil@cmu.edu)

ABSTRACT

This paper describes a smart radio network that can facilitate communication of radios operating under congested conditions in the 2.4 GHz unlicensed band. The network detects legacy systems spectrally, temporally, and spatially, and implements a novel avoidance technique to reduce or remove interference with the legacy systems. The system was designed by a team of students from Carnegie Mellon University for the 2008 Smart Radio Challenge sponsored by the Software Defined Radio Forum.

1. INTRODUCTION

There has been an increasing interest in the study and research of cognitive radio systems for various reasons. One reason in particular is spectrum efficiency [1]. In the 2008 Smart Radio Competition, the Software Defined Radio (SDR) Forum gave a problem scenario where cellular networks operating in a particular portion of spectrum have reached maximum capacity, causing degradation in the quality of service (QoS). Participating teams were charged with the task of resolving this problem through mapping the spectrum and finding unused and under-used spectrum then allocate these available spectrum resources to users in their secondary network [2]. The student team from Carnegie Mellon University, Team Plaid, designed a system for this problem scenario. This paper describes and gives an overview of the system designed by Team Plaid.

The spectrum usage can be mapped temporally and spatially. Temporal spectrum mapping involves associating time information with the spectrum sensed by the radios or network. In spatial spectrum mapping, the data is associated with physical locations. In the system designed by Team Plaid, the network was able to determine the times and locations for which certain parts of the 2.4 GHz unlicensed band were being used. The system was then able to use the unoccupied spectrum for secondary communications.

2. DESIGN SOLUTIONS

The radio network designed by Team Plaid consists of three types of radios: the base station, user radios that desire to establish a communication link with other users in the network (e.g. voice, data, or text), and radios that provide sensing functions only. The base station serves as a central hub that coordinates sensing of the band and allocates communication channels for the users. To minimize storage capacity needed by the user, the base station is also responsible for creating, storing, and maintaining a database of the analyzed data and the information learned. The user radios communicate on the channels identified by the base station, and immediately stop transmitting when a primary user is sensed. The remaining radios are responsible for collecting data about the entire spectrum and transmitting it to the base station via the allocated channels.

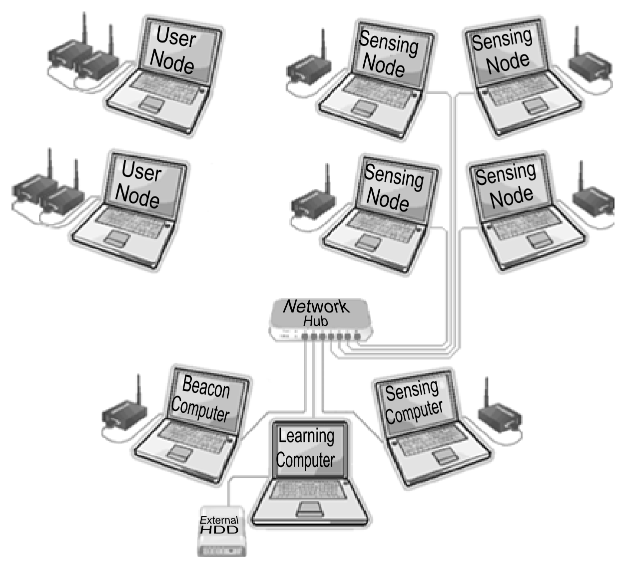

Fig. 1. Diagram of the complete radio network designed by Team Plaid.

The four sensor nodes shown in Figure 1 constantly monitor the spectrum as does the sensing computer at the base station. They dump all of their sensing data onto the external hard disk drive. The learning computer then processes that data to map the spectrum. After a period of time, the learning computer then allocates spectrum to the user nodes. The beacon computer then transmits a beacon signal signifying what channels are available for the user nodes. Once interference on the beacon channel is detected, the beacon computer moves to a different channel to avoid the primary users.

3. HARDWARE

The radio network consists of Universal Software Radio Peripherals (USRPs) from Ettus Research. The user nodes run the USRP with the standard hardware. For the sensor nodes, the USRPs were modified to enhance the monitoring bandwidth. The base station also uses a modified USRP for sensing as well as a wideband signal conversion module (SCM). The SCM achieves better sensing performance than what is possible for the USRPs [3]. This section describes the modifications made to the USRP code to improve sensing performance. The general schematic of the SCM is also detailed along with key performance parameters.

3.1 Modified USRP

The USRP has limited bandwidth due to the maximum communication speed between the radio and the host computer. The maximum bandwidth that these radios can sense is 8 MHz [4]. To sense the entire spectrum the channel on the USRP must be changed over 10 times. Each channel change introduces a significant delay of ~400 ms. By sensing a wider bandwidth, the time required to sense the entire ISM band can be reduced.

The ADC on the USRP samples complex values at 64 Msps. Therefore the maximum possible bandwidth that the USRP can receive is 64 MHz. The samples can only be read out of the USRP at 8 Msps, so the data sampled at 64 Msps must be stored in memory until it can be sent to the host. When the memory is nearly empty another block of 64 MHz data is sampled and stored in memory. This buffering procedure is illustrated in Figure 2 [4].

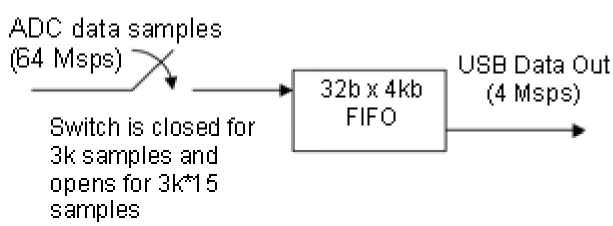

Fig. 2. Schematic of the 64 MSP buffering system

The length of the data blocks determines the frequency resolution after performing a fast Fourier transform. The length of these blocks is limited by the memory on the USRP FPGA. The FPGA is a Cyclone I from Altera with 234kb of RAM. Each complex sample is 32 bits long so the maximum buffer size is about 7.3×2^{10} samples. The memory on the FPGA must be shared with other functions such as the transmit buffer. The memory left over for the receiver FIFO is about 4×2^{10} samples. The 64 MHz data block size is 3×2^{10} to prevent the receiver FIFO from overflowing. This block size results in a minimum frequency resolution of 20.8 kHz [5].

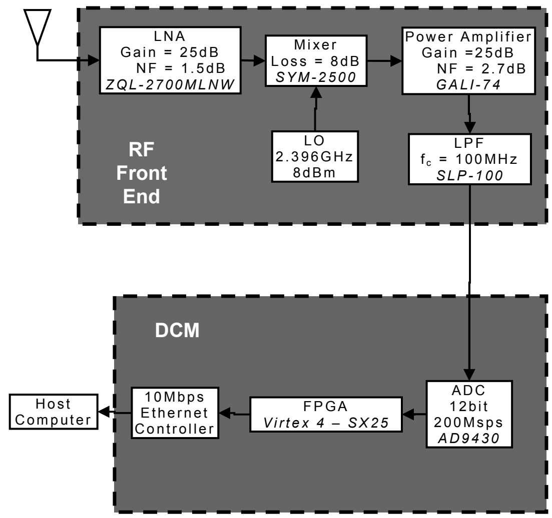

Fig. 3. Data path in the SCM, All parts are from Minicircuits except the ADC and FPGA which are from Analog Devices and Xilinx respectively.

3.2 SCM

The signal conversion module consists of two parts: the RF front end and the digital conversion module (DCM). The RF front end mixes and filters the entire ISM band to baseband. The DCM samples the base band data, stores the samples and sends the stored data to a computer. The SCM basically performs the same operation as the modified USRP but with a wider bandwidth and larger memory [3] [6].

The RF front end is shown in Figure 3. The RF signal received by the antenna is amplified by a low noise amplifier. The amplified signal is then mixed down with an LO of 2.396 GHz. The baseband signal is then amplified again and filtered by a 100 MHz anti-aliasing filter.

The DCM samples the baseband signal from the RF front end at 200 Msps. The ADC on the DCM gives 12 bit real samples, not complex samples. The FPGA records blocks of the sampled data in the same manner as the modified USRP, but the amount of memory on the DCM is much larger than the USRP. The RAM on the DCM FPGA can hold 174,763 samples. which gives a frequency resolution of 1.14 kHz [6] [7].

The samples on the FPGA are sent to a host computer over 10 Mbps Ethernet. The actual data rate is about 1.7 Mbps, which converts to a data block every 1.2 s. The limiting factor in the data rate is due to latency in the reception code at the host.

The dynamic range of the system was measured with a narrowband signal at 2.45 GHz inserted at the RF front end input. When the signal strength was above -40 dB clipping in the ADC occurred. The signal was below the noise floor when the power was below -107 dBm. Therefore, the receiver has a dynamic range of about 67 dB. There is no automatic gain control in the radio so high dynamic range is necessary to detect a wide range of signals.

4. SOFTWARE

All computers connected to a radio, excluding the sensing computer at the base station, are operating in Windows with Team Plaid's Pipeline system to take advantage of the benefits of MATLAB. The Pipeline system allows users to use the USRP in a Windows OS. The USRP is commonly used with Linux. Pipeline allows the operation in Windows. It does exactly what its name states. It pipelines data from one OS to another through a virtual machine (VM). A VM was setup to run a Linux OS and share files with the host computer. Using Python and GNURadio, the Linux guest computer grabs data from the USRP and saves it to a binary file in the share folder. This allows the Windows host computer to use these files. The majority of programs were developed in MATLAB, excluding the code to transfer data from one OS to another.

5. SENSING

The idea behind the sensing stage is to build an accurate database of all energy present in the 2.4 GHz Instrumentation, Scientific, and Medical (ISM) band through sensing in the temporal and spatial domain [8].

When the term sensing is used in this paper it refers to a fast Fourier transform (FFT) of the data collected by the radios. The frequency bins of the FFT correspond to the center frequencies and channel bandwidths of the network.

The FFT data from all of the radios are collected and stored in a database. The resolution bandwidth of this database is based on the smallest channel bandwidth that the base station can assign. The secondary network is structured to operate with five different channel bandwidths: 1.25 MHz, 1 MHz, 625 kHz, 312.5 kHz, or 62.5 kHz. Each channel has a center frequency, $f_{B,n}$, that is determined by

$$f_{B,n} = 2399.5 + nB \text{ MHz}, \quad (1)$$

where B is the channel bandwidth, n is the channel number, and $nB \leq 2500.6875$ MHz. For example, if operating at a resolution bandwidth of 312.5 kHz, the frequency of the 18th channel is 2405.125 MHz. The center frequencies are based on the channel and sub channel frequencies of legacy systems such as 802.11b/g to maximize sensing accuracy.

The sensing process starts with the collection of data from the sensing nodes. The sensing nodes capture three data blocks with center frequencies between 2.4 and 2.5 GHz as shown in Figure 4.

The anti-aliasing filters on the USRPs do not completely filter out everything outside of the 64 MHz window. This causes higher attenuation to signals near the edges of the 64 MHz window. To ensure a reasonably uniform gain response across the band and to avoid aliasing, the windows are overlapped by 32 MHz, and only the center 32 MHz of each channel is used. The 64 MHz windows are centered on the frequencies of 2417 MHz, 2450 MHz, and 2483 MHz. The data has a frequency resolution of 62.5 kHz. After the sensor nodes obtain the data it checks to see if the base station is accepting data. If the base station is accepting data, all of the data between frequencies 2399.5 MHz and 2500.6875 MHz are then sent to the base station for processing. Otherwise the data is discarded. The sensor nodes continuously send this information to the base station upon request.

The sensing computer at the base station behaves in a similar manner but its hardware is better suited to detect signals with bandwidths less than 2 MHz and it has the bandwidth to capture the entire band in one data block. In the preliminary stage, only the sensing data from the base station is available for processing. After the preliminary sensing, the base station will only be responsible for collecting data (i.e. magnitude and phase) in the channels used for coordinating the sensing and communications. However, in this design we coordinate the transfer of sensing data over a separate backbone network via Ethernet cable. Therefore, the base station will also be used to sense the entire spectrum as well.

The user nodes are also responsible for sensing on their allocated channels. The collected data can then be transmitted back to the base station via the same channel, or via the wired backbone network [9].

All of the sensing data is stored on the external hard disk drive for processing by the learning computer at the base station.

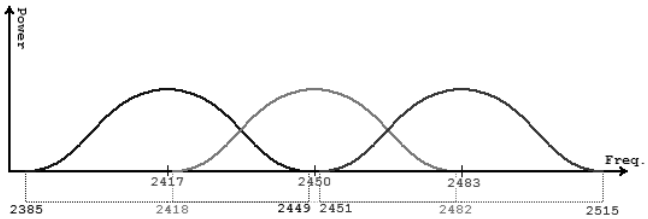

Fig. 4. Spectrum windows for the USRPs

6. LEARNING AND INTERFERENCE AVOIDANCE

The process of analyzing the data is called the learning stage [8]. In the learning stage the learning computer at the base station performs various calculations on the sensing data. The learning computer estimates the frequency, energy, type of signal, transmission time, and relative location. This information will make up the database. From this database the base station can determine the channels in which user-to-user communication can be facilitated.

6.1 Energy Detection

Using predetermined noise floor calculations for the radios, the learning computer searches for energy that exceeds the noise floor by holding peak energy values from the data obtained. The sensor nodes detect signals with bandwidths larger than 2 MHz, while the more capable radio attached to the sensing computer detects signals less than 2 MHz as well as digitizes the entire ISM band at once. After analyzing all the data, a list is populated containing the spectral location of the energy. The learning computer then attempts to classify detected signals based on the measured bandwidth [10] [11] [12]. For example, if the signal occupies approximately 20 MHz and is centered on one of the 802.11 frequencies, the system classifies the signal as an 802.11 transmission. The energy detection also aids in the estimation of the bandwidth and center frequency for all other signals. This is beneficial in classifying other devices [10] [11].

6.2 Temporal Detection

After determining where the energy is located in the spectrum, the learning computer measures the duration of the transmission based on the data from the sensing computer. The channel duration assists in determining whether or not the signal is frequency hopping [12].

6.3 Location Detection

The ratios of the received signal power among the sensing nodes are used to estimate the spatial location of the primary transmitters. Given that the sensing radios are similar, meaning the same type of antenna and hardware, the received power specifies the radius of a circle that defines the location of the transmitting radio. With three nodes three unique circles can be generated and their intersection signifies the location of the transmitting radio. Although this scheme should work well in an open field environment, for indoor locations we found that range errors caused by multipath often prevented an accurate position estimate [13].

6.4 Duty Cycle

The learning computer also calculates the duty cycle of the spectrum. Any spectrum with a duty cycle below a predetermined threshold is identified as under used or unused spectrum and assigned to the secondary users. Other papers have used this method to determine the occupancy of the spectrum [14].

6.5 Interference Avoidance

Once a channel is assigned, the beacon computer begins to transmit a beacon signal on the channel. The beacon serves a dual purpose: it identifies the channel assignment to the users, and provides a means for interference detection. The key concept underlying our approach is the use of a beacon signal that is very sensitive to interference.

The beacon used in our implementation is a multicarrier signal with a binary code determining whether a particular carrier is on or off. This provides a unique spectral signal to enable user nodes to unambiguously distinguish between beacons and other user signals. The beacon uses up to twenty carriers with a spacing of 62.5 kHz for a maximum bandwidth of 1.25 MHz. A threshold is applied to the spectrum amplitude to create a binary signal for correlation with the beacon code. The presence of an interfering signal causes errors, resulting in a reduction in the magnitude of the cross-correlation peak. Rather than setting the decision threshold to half the peak value, the threshold is set close to the "zero" symbol level to maximize the sensitivity to interference. This way any slight increase in spectral power will create an error. This is equivalent to monitoring the frequencies where the beacon spectrum is zero for the presence of interfering signals. If excessive errors are observed, the beacon is immediately turned off. When the user nodes see that the beacon has been turned off, or they too detect errors on the beacon signal, they also immediately stop transmitting.

An alternative scheme with more efficient use of the spectrum could be implemented using a single carrier. For example, a low data rate signal using on-off-keying (OOK) modulated with a PN sequence could be used. As in a conventional communication system, the baseband OOK signal could be converted to a digital signal by comparison with a threshold. With the multicarrier implementation, the threshold could be set very near to the "zero" symbol level for maximum sensitivity to interference. This way, any slight interference appearing when the OOK signal is off would introduce an error, and the number of errors would be an indication of the presence and severity of interference. With a low chip/data rate, the bandwidth required by the beacon would be minimal.

However, the temporal synchronization necessary for processing a time-domain signal as described above presents considerable challenges with the hardware available for our system. Consequently, we chose to implement the analogous frequency-domain version of the scheme. Although the frequency-domain version uses more spectrum than what is required for the beacon function and therefore is not the optimum implementation for a congested spectrum, it does allow the principle to be illustrated without the requirement of tight temporal synchronization.

7. USER NODES AND APPLICATIONS

The user nodes constantly sweep the spectrum looking for the beacon signal. Once the beacon signal is detected, the user nodes lock onto the beacon signal. The user nodes then set up a secondary network based on 3.75 MHz channels with the beacon occupying the center 1.25 MHz. On either side of the beacon is 1.25 MHz divided into twenty 62.5 kHz sub-channels. The first user selects and occupies the number of sub-channels required for the bandwidth of the desired transmission. Subsequent users may select any remaining unused sub-channels. Once the user nodes are locked on they can begin communication. In the demonstration system, applications were implemented on the user nodes for sending text, voice, and photos. The user nodes continue to transmit as long as they are able to detect the beacon signal.

The first application sends an ASCII text string to another user node specified by the user ID, whenever there is an available adjacent channel to the beacon. The second application is voice messaging. This application records a 10-second voicemail locally, and sends the voice message to another user node. The third application is "real-time" picture sharing. There is a web camera installed on the user node that captures pictures periodically (for example, every 5 seconds). Once the photo sharing application on the user node is started, each user is able to send the latest capture to the other user.

Each user node uses MSK modulation. For the MAC layer, a predetermined channel setup based on the beacon signal is used, or they can deploy carrier sensing to determine what channel to use. These are the only two network layers the radios consider.

8. SUMMARY AND CONCLUSION

In conclusion, the operation of the system can be summarized as follows.

Consider the case where Team Plaid's system is placed in a room with an 802.11b/g access point (AP). Initially, the learning computer at the base station will begin to search for open spectrum using its local measurements as well as measurements obtained from the sensor nodes. It will spectrally and spatially find and locate the 802.11 AP using the relative signal strengths from the various sensing nodes. After approximately 30 seconds the base station will select an unoccupied channel for the user nodes. The base station will then begin transmitting a beacon signal in the center of an available 3.75 MHz segment of spectrum. The user nodes detect the presence of the beacon using cross-correlation with the beacon code, and thus identify the available channel. The first user chooses and occupies the number of adjacent 62.5 kHz sub-channels required for the bandwidth of the transmission. Subsequent users select from among any remaining unused sub-channels. Throughout the entire process the sensor nodes and sensing computer are constantly monitoring the spectrum and sending information back to the learning computer. The learning computer will constantly check for the primary users in the spectrum as well as the presence of the beacon signal.

If the AP changes its frequency to the channel that the user nodes are occupying, the user nodes will be able to detect it by the increase in beacon errors. The user nodes will then stop transmitting, and begin searching for a new beacon signal. Similarly, when the learning computer detects excessive beacon errors from either the sensing nodes or the sensing computer, it will turn the beacon off. It will then consult its most recent space-time spectral map and select a new available channel. The beacon computer will then transmit a beacon signal in the center of the new channel. When the new beacon is detected by the user nodes, they will reset their frequencies to sub-bands adjacent to the new beacon as before. In this manner, the network will constantly avoid the channels that are occupied by the primary users.

REFERENCES

[1] B. Fette, *Cognitive Radio Technology*. Burlington, MA: Newnes, 2006.

[2] S. Forum, "Smart Radio Challenge 2008: Sample Problems (http://www.radiochallenge.org/08SampleProblems.html)," , Phenoix, AZ, 2008.

[3] Departments of Computer Science and Electrical and Computer Engineering Carnegie Mellon University, "A Controlled Wireless Networking Testbed based on a Wireless Signal Propagation Emulator

(http://www.cs.cmu.edu/~emulator),". Pittsburgh, PA, 2007.

[4] Eric Blossom, "Exploring GNU Radio (http://www.gnu.org/software/gnuradio/doc/exploring-gnuradio.html), November 2004

[5] Altera, Inc., "Cyclone FPGA Family Data Sheet (http://www.altera.com/literature/hb/cyc/cyc_c5v1_01.pdf)," 2007

[6] Borries, K.C.; Judd, G.; Stancil, D.D.; Steenkiste, P., "FPGA-Based Channel Simulator for a Wireless Network Emulator," Vehicular Technology Conference, 2009. VTC Spring 2009. IEEE 69th , vol., no., pp.1-5, 26-29 April 2009

[7] Xilinx, Inc., "Xilinx Digital Signal Processing : XtremeDSP Development Kit for Virtex-II Pro (http://www.xilinx.com/xlnx/xebiz/designResources)," 2007

[8] J. Mitola III, "Cognitive radio: An integrated agent architecture for software defined radio," Doctor of Technology dissertation, Royal Inst. Technol. (KTH), Stockholm, Sweden, 2000.

[9] IEEE 802.11n Committee, IEEE Draft Standard for Information Technology-Telecommunications and information exchange between systems-Local and metropolitan area networks-Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications - Amendment 3: 3650–3700 MHz Operation in USA. New York, NY: IEEE, 2007.

[10] 802.11-2007 IEEE Standard for Information technology-Telecommunications and information exchange between systems-Local and metropolitan area networks-Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications

[11] Yucek, T.; Arslan, H., "Spectrum Characterization for Opportunistic Cognitive Radio Systems," Military Communications Conference, 2006. MILCOM 2006. IEEE , vol., no., pp.1-6, 23-25 Oct. 2006

[12] M. Gandetto, M. Guainazzo and C. S. Regazzoni, "Use of Time-Frequency Analysis and Neural Networks for Mode Identification in a Wireless Software-Defined Radio Approach", Eurasip JASP, Special Issue on Non Linear Signal Processing and Image Processing, Vol. 13, pp. 1778-1790, Oct. 2004.

[13] A.K.M. Mahtab Hossain, Hien Nguyen Van, Yunye Jin and Wee-Seng Soh, "Indoor Localization using Multiple Wireless Technologies," IEEE MASS, Pisa, Italy, Oct. 2007.

[14] Vilmpoc, Max, "Dupont Circle Spectrum Utilization During Peak Hours: A Collaborative Effort of the New America Foundation and The Shared Spectrum Company", New America Foundation, 2007