

# Virginia Tech CWT's Smart Radios – Challenges, Solutions and Lessons from the 2007 & 2008 Competitions

Alexander R. Young, Mark D. Silvius, Terry Brisebois, Feng Ge, Rohit Rangnekar, Charles W. Bostian  
Center for Wireless Telecommunications, Virginia Tech, Blacksburg, VA  
Email: {alex.young, msilvius, tbrisebo, gef, rangnek, bostian}@vt.edu

**Abstract**—This paper presents the Wireless@Virginia Tech CWT entries in SDR Forum's Smart Radio Challenge. The SDR Forum's Smart Radio Challenge offers student teams the opportunity to design and implement actual radio systems that address issues presented by real world problems to public safety first responders. In 2007, the Smart Radio Challenge problem chosen by CWT concerned recognizing and avoiding primary users while maintaining a communication link with a user-specified minimum QoS. In 2008 we extended the previous year's solution to a network that automatically provided an emergency response communications system in the event of damaged or destroyed infrastructure. Both design experiences contained risks that yielded lessons to be learned. Factors like hardware, software, scenario testing, and design vs. implementation trade offs all had to be weighed. For our 2007 and 2008 entries, this paper presents some of those considerations and choices made, conclusions and lessons learned, and how our experiences might guide future efforts, both in this year's challenge and in other projects.

**Index Terms**—smart radio challenge, software define radio

## I. INTRODUCTION

Since 2001, the Wireless@Virginia Tech Center for Wireless Telecommunications (CWT) has developed solutions to the problems associated with disaster response and public safety communications, from the first, applying cognitive radio (CR) and software defined radio (SDR) [1]–[4]. *Smart Radio* is a term used to describe a general class of devices that includes software defined radio (SDR), frequency agile radios, cognitive radios, and associated technologies [5]–[9]. For the purpose of the SDR Forum Smart Radio challenge, we define a Smart Radio (SR) as a “sensing radio, programmed to respond to changes in its environment in an innovative way,” [10].

The SDR Forum's Smart Radio Challenge [11] offered Virginia Tech students the opportunity to apply research experience in solving practical communications

problems and developing and demonstrating real world prototypes.

This paper presents CWT's 2007 and 2008 entries in the Smart Radio Challenge. We review each year's challenge scenario and design problem, and then present our smart radio. We give an overview of the entire system, highlighting its key components and major operational modes. A review of the team's process discusses risks, stumbling blocks, lessons learned, and contributions.

The remainder of this paper is organized as follows. Section 2 covers the 2007 smart radio challenge, discussing first the challenge problem itself, then presenting Virginia Tech's solution to the problem, finally reviewing and reflecting on issues, problems and lessons learned. Section 3 covers the 2008 challenge, again presenting first the problem itself, Virginia Tech's solution, and reflections. Section 4 discusses future work and the manner in which we might apply the lessons we have learned. Section 5 presents additional thoughts on the Smart Radio Challenge as a whole.

## II. 2007 CHALLENGE

This section presents the 2007 smart radio challenge, first discussing the challenge problem itself, and then presenting the solution developed by CWT. This section concludes with a risk review.

### A. Problem

2007 was the inaugural year for the SDR Forum's Smart Radio Challenge. Teams were offered a selection of scenarios from which to choose. We chose to address the problem of spectrum access for first responders. A large city experiences a major crippling earthquake. As a result, all communications infrastructure is destroyed. In this scenario, emergency first responders must establish a makeshift command center and establish a temporary communications infrastructure. Teams addressing this

scenario were tasked with developing a system that fulfilled four key points [12].

- 1) Find available spectrum within a predefined band
- 2) Rendezvous with an intended receiver
- 3) Transmit data over that band with a pre-determined Quality of Service (QoS) in urban conditions
- 4) Guarantee non-interference with primary users.

### B. Solution

We generated a list of goals to guide system development. In general the system should be mobile and user friendly. Operationally, the system should:

- Sense and classify the RF environment
- Interoperate with other smart and family radio service (FRS) radios while reducing mutual interference
- Operate in Master-Slave or Infrastructure modes
- Adhere to FCC spectrum regulations

Our solution leveraged work previously developed by CWT, specifically the Public Safety Cognitive Radio (PSCR) [13]. Using the functionality provided by the PSCR, we developed the 2007 Smart Radio (SR-1) with a modular system that exchanged data between components using an eXtensible Markup Language (XML) based messaging scheme. The components that make up the architecture of SR-1, as shown in Figure 1, are the Master Control, User Interface, Sensor, FRS and knowledge database, and Radio Framework. The modules are connected to each other through the use of network sockets, and information is exchanged between modules in a standardized format using XML.

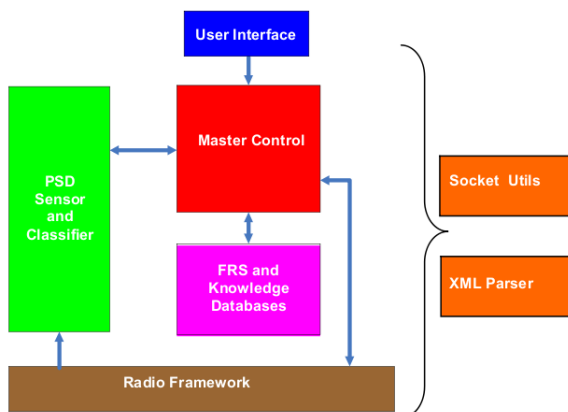


Fig. 1: 2007 Smart Radio Architecture [9].

1) *Master Control*: The Master Control is the brain of SR-1. It coordinates and controls the functions of the entire radio. Through the user interface, the MC

receives commands from the operator, and provides visual feedback to the operator. Upon receiving operational commands, the MC uses one or more of the other modules to execute/fulfill the operators wishes. The MC will pass a run command to the environmental sensor, causing the sensor to sweep the RF spectrum looking for energy. The sensor may detect users and classify their signals; this information is returned to the MC, which stores it in a knowledge database. While the sensor provides the MC with spectral awareness, the knowledge database facilitates signal memory. Using the information in the knowledge database as well as reference information in the FRS database, the MC can direct the radio framework to configure and connect to any of the identified signals as desired. The MC uses two protocols to manage dynamic spectrum access (DSA) in a crowded RF environment. The protocols are called Channel Change Protocol (CCP) and Rendezvous Protocol (RP) [9].

SR-1 operates in an environment where it is often considered a secondary user; i.e., SR-1 often communicates using spectrum hole channels of opportunity, which it is allowed to do so long as it does not cause interference to primary users. SR-1 must vacate a channel when the primary user returns to use the channel. SR-1 must then choose another unused channel and reestablish communications. RP allows SR-1 to quickly locate another smart radio in a given frequency band. RP is implemented using a series of frequency-hopped beacons containing node specific information. Similar to Bluetooth's inquiry messages, RP transmits a beacon on a particular frequency, listens for a response, and then moves on to another predetermined frequency. Rendezvous occurs when two smart radios, each using RP, connect. Specifically, rendezvous occurs when one smart radio transmits its beacon on a particular frequency, and another smart radio is listening for a beacon on the same frequency and receives the transmission [10].

CCP is a spectrum sharing protocol that allows primary and secondary users to coexist in the same frequency band without interference. As implemented in the CWT smart radio, CCP detects the presence of an analog FRS primary user in a previously unoccupied channel. Upon detecting the primary user, CCP causes the smart radio to reconfigure itself and move to an alternate set of channels, known as fallback channels. If the destination channel is occupied, as determined by the primary user detection routine, the radio continues to proceed through a series of predetermined fallback channels, until the radio finds one that is vacant. At this time, the smart radio resumes interrupted communications [10].

2) *User Interface*: The user interface (UI) is the module that interfaces directly with the operator, a public safety officer. Using the UI, the operator can scan and classify the environment, looking for active communication signals or rendezvous beacons; transmit rendezvous beacons using RP; or initiate communication with an FRS or smart radio using a variety of waveforms [14].

3) *Sensor*: The smart radio's intelligence comes in part from its situational and environmental awareness. This awareness is provided by the environmental sensor, in conjunction with a knowledge storage database. The sensor performs signal detection and classification on any signal in the desired frequency band. Signal detection is implemented using a Fast Fourier Transform (FFT)-based energy detector. The PSD sensor compares collected samples to a predetermined threshold. Anything that registers above the threshold is noted and relayed to the signal classifier. The signal classifier is based on a K-nearest neighbor (KNN) algorithm, classifying signal samples based on their time domain characteristics. Different public safety waveforms have different characteristics when compared at the same signal to noise ratio (SNR). In this manner, the sensor is able to distinguish between analog and digital waveforms, classifying FM and M-PSK waveforms [13].

4) *FRS and Knowledge Databases*: The SR-1 combines information about standard FRS waveforms and knowledge about the observed RF environment in its memory. The FRS database is used to store FCC-approved waveform definitions for FRS radios. These waveform definitions are used when the smart radio needs to communicate with legacy FRS radios. The knowledge database maintains information about the RF environment as provided by the environment sensor. The MC uses this spectral information to determine whitespace in which it can safely establish or resume communications [9].

5) *Radio Framework*: The radio framework provides SR-1's implementation of a transceiver. The framework provides all the digital signal processing (DSP), filtering, sampling, modulation and similar functionality that allows radios to be implemented in software. Radio functionality is provided by the GNU Radio software [15] and the Universal Software Radio Peripheral (USRP) [16] RF front end. Radio flowgraphs, as shown in Figure 2, are implemented in C++ and Python and provide transmission of voice and data using analog or digital modulations. A rudimentary carrier sensing medium access control (MAC) is implemented using GNU Radio, allowing user detection and shared multiple access.




Fig. 2: Transmit and receive radio flowgraphs [9].

### C. Risk Review

CWT has a history of using GNU Radio in its radio development. Our team made use of previous students' expertise to get up to speed quickly in software radio development. In addition to familiarity with GNU Radio, CWT students have long used the USRP as the radio hardware front end. We moved forward, developing software with GNU Radio based on the USRP hardware. Using the sponsor-provided Lyrtech Small Form-Factor (SFF) SDR development platform [17] would have required a great deal of embedded systems programming, which we lacked. Furthermore, a very large commitment of resources would have been needed to recreate software, modules, functions and algorithms on a new platform. In the end, the SFF arrived on-site late in the development cycle; our decision to develop on an alternate platform was the right decision for several reasons.

In addition to leveraging CWT knowledge on alternate SDR hardware and software platforms, we decided to build its smart radio on the PSCR, a product previously developed by CWT. The PSCR is an implementation of cognitive radio focused on the public safety domain. The PSCR allows emergency response and public safety officials to scan the public safety radio spectrum, identify and classify any non-proprietary standards-based public safety waveform, and reconfigure itself to communicate with any of the identified signals as desired. The PSCR also has the ability to act as a repeater or bridge two otherwise incompatible waveforms. More information on the PSCR can be found in [13], [18].

## III. 2008 SMART RADIO CHALLENGE

This section presents the 2008 Smart Radio Challenge, first discussing the challenge problem itself, and then presenting the solution developed by CWT. This section concludes with a risk review.

### A. Problem

The 2008 problem involved a team of first responders on the scene following a major incident in a city subway. Emergency response teams are equipped with a radio that is capable of two-way voice operation. However, any communications infrastructure available in the subway has been destroyed. Participating challenge teams were responsible for developing a smart radio system that would automatically create an ad-hoc extension to the existing communications network, allowing voice communications to be relayed into and out of the subway incident area and connecting to infrastructure unaffected by the subway incident [19].

### B. Solution

In designing a smart radio solution to address all the aspects in the 2008 challenge, we used the following design guidelines:

- The smart radio will automatically create an ad-hoc extension to an existing communications network such that voice communications can be relayed to (from) the incident site from (to) the above communications infrastructure.
- The network extension will utilize peer-to-peer links among the radios and reconfigure at least one radio as a repeater.

In addition to the above design requirements, the smart radio incorporates and improves upon features introduced in the earlier competition. With respect to topology, the public safety ad-hoc network should incorporate at least six smart radio nodes. One node is a hidden node, representing a public safety official who is completely cut off from communications infrastructure. Three nodes act as repeater radios, which form the ad-hoc network with the tunnel and relay voice and data from the hidden node to the tunnel's entrance. A command node, situated at the entrance of the tunnel, acts as a gateway bridging the tunnel network to the intact above-ground communications infrastructure. One radio node represents the above-ground infrastructure.

Our 2008 Smart Radio built upon the successes of the previous year. The 2007 Smart Radio (SR-1) served as the base platform for the 2008 Smart Radio (SR-2). In addition to all the features provided by the base platform and discussed in Section 2, SR-2 incorporated numerous additions and technical advances. The SR-2 architecture was similar to that used in 2007, as shown in Figure 3. However, every component gained functionality.

1) *Master Control*: As in SR-1, the Master Control (MC) remains the brain of the system. The MC retains all its capabilities from SR-1, but adds support for additional

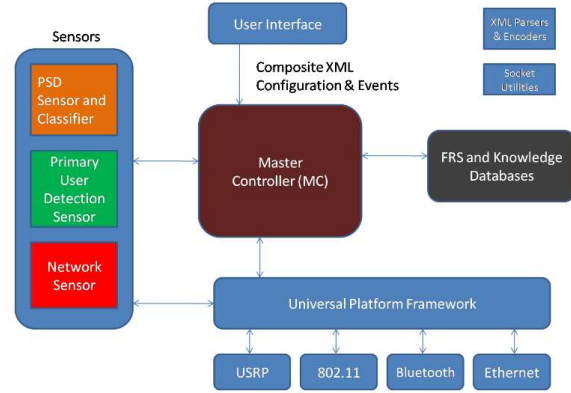


Fig. 3: 2008 Smart Radio Architecture [20]

sensors and additional radios interfaces. A significant addition to SR-2 is the implementation of event-driven control. Trigger events and corresponding responses are defined by the operator using the user interface. If the MC observes the trigger event, the MC will cause SR-2 to reconfigure itself in the manner previously indicated by the operator. This event-driven functionality allows SR-2 to rapidly change modes from standard radio node to ad-hoc relaying network bridging the remote hidden node to the uncompromised communications infrastructure.

The event-driven functionality would be triggered by a loss of Internet connection. Upon this event, the MC would search for an alternate way to connect to the Internet. It would try a wired Ethernet connection first, followed by a Wi-Fi connection. If neither was available, it would go into ad-hoc mode. In ad-hoc mode, it would search for another user, via Bluetooth, with an Internet connection. Upon finding one, it would connect to that user. At any time, when the MC possesses an Internet connection, it will allow others to connect to it, and will share its Internet connection with them, and bridge the network appropriately. This event-driven logical flow ensured that a Bluetooth-based minimum-spanning tree architecture would be created to share any available Internet connectivity and to rebuild the data network.

2) *User Interface*: SR-2's graphical user interface (GUI) is the critical connection between the radio operator and the radio itself. With one touch, the operator can scan the environment using multiple radio interfaces, and see any available Ethernet, WiFi, Bluetooth and public safety legacy radio connections all on one screen. The GUI also provides an interface through which the operator can define RF trigger events, such as a lost connection, and the processes that the radio should



automatically execute in response to those triggers.

3) *Sensor*: SR-2's sensor subsystem retained the PSD sensor from 2007 as the primary observation mechanism of the wireless environment, and added a network sensor. The network sensor provided information on the status of network connectivity. If the network connection status changed, the network sensor would notify the MC. The MC would respond to this trigger event in a manner previously determined by the operator, and SR-2 would reconfigure itself to adapt to the network change, connecting to other networks as required.

4) *Knowledge Databases*: Commensurate with SR-2's increased capabilities, the knowledge databases were increased to keep track of additional data, including connection information for the Ethernet, WiFi, Bluetooth and legacy public safety radio domains. This increased functionality allows SR-2 to bridge not just public safety radio communications, but network communications as well. For example, it allows a voice-over IP digital voice stream to be sent to an analog public safety radio.

5) *Radio Framework*: The radio framework used by SR-2 received significant upgrades over 2007 and SR-1. SR-2 maintained its legacy radio compatibility, as indicated by the use of the USRP. The smart radio-centric and legacy radio-centric framework of 2007's SR-1 was replaced with a universal framework, supporting multiple radio interfaces. In addition to the traditional SDR functionality, the universal framework supported commercial off-the-shelf interfaces including 802.11WiFi, Bluetooth, and Ethernet. Thus, SR-2 could communicate over any of these interfaces, using both wired and wireless media. The universal framework also supported bridging any of the radio interfaces to any other. Efficient use of the various communication methods was accomplished through the incorporation of On-Demand Link State Routing (OLSR) and a custom Bluetooth scatternet formation and routing protocol.

### C. Risk Review

In developing the CWT Smart Radio 2008, the team ultimately decided to put a larger emphasis on the use of commercial off-the-shelf (COTS) hardware devices. The reason for this was twofold. First, the COTS devices readily met the requirements set forth in the competition problem statements. For example, the requirement for mobile radio nodes to establish a multi-hop ad-hoc network in the 2.4 GHz band was already fulfilled by the 802.11 WiFi standard and open source routing protocols like OLSR. Second, existing SDR front-end technology such as the USRP could not deliver the QoS performance required at the reliability level desired when compared to 802.11 WiFi.

The end result of this decision was that CWT had to meet the flexibility requirements of a SDR platform by running and switching between multiple fixed-waveform COTS front ends connected in parallel, rather than running a single reconfigurable custom SDR front end, which could dynamically reconfigure its transmitted waveforms. We considered either strategy a valid solution, but chose the COTS method as a quicker, less expensive, and more reliable solution while still fulfilling the requirements set forth in the competition.

The choice to use COTS components had unexpected consequences. Using COTS devices turned the implementation of SR-2 research and development effort into primarily an integration effort. Systems integration is an important area in engineering industry, but one not often encountered in academic research. CWT found that the integration problem was not trivial. Learning the COTS device APIs and modifying the devices' drivers for integration into the radio framework required significant effort.

While use of COTS devices in a SDR competition, a field accustomed to using non-traditional and SDR-specific hardware, might be considered outside the scope of the competition, CWT has a different perspective. COTS devices are full radio systems in their own right. They have very specific tasks and are very well suited to performing those tasks. As a result, there is little effort expended in turning such devices towards SDR goals. CWT, however, did just that. SR-2 fully integrates COTS devices into its radio framework along side traditional SDR hardware, and SR-2 has the ability to control them all, switching seamlessly between them.

## IV. GOING FORWARD

The lessons and technology from two years of SDR Forum Smart Radio Challenges will be useful in two major ways, at CWT and in the 2009 challenge.

### A. CWT

The radios, SR-1 and SR-2, are fully developed software radios with several distinct and unique capabilities. These radios can serve as development platforms for future CWT efforts. SR-1 and SR-2 served as showcases for DSA technology such as CCP and RP, multi-interface routing and bridging, and event-driven automatic radio configuration. These technologies, and SR-1 and SR-2 themselves, can find a place in future CWT research. CWT's current research efforts include investigation of heterogeneous ad-hoc networks, physical layer link optimization, and distributed sensing.

## B. 2009 Challenge

In previous years, we successfully leveraged existing technology developed by current and former members of CWT. This tradition will continue in the 2009 challenge. The challenge this year presents a scenario where there has been a catastrophic natural disaster and all communications infrastructure has been eliminated. Emergency first responders are arriving on the scene and facing interoperability and mutual interference problems. Challenge teams are tasked with developing a cooperative sensing system that creates and maintains a database of public safety radio nodes, tracking node location, waveform physical layer characteristics, and public safety team associated with each node. This challenge problem is sufficiently unique that the Virginia Tech team this year will not build its solution on either the SR-1 or SR-2 platform. We will however leverage existing technology developed in CWT. CWT has been developing distributed sensor networks, advanced signal detection methods, DSA management techniques and temporary emergency response communications infrastructure [21]–[25]. These and more will serve as the basis for Virginia Tech's 2009 challenge entry.

## V. ADDITIONAL THOUGHTS

The Smart Radio Challenge continues to be a unique forum for smart software radio design and exhibition. The opportunity to meet other student teams is always exciting, and we look forward to future opportunities to interact with these fellow researchers. There has been little such opportunity in the past, and we hope that future Challenges will allow teams the chance to interact and observe each other's projects. Finally, the Smart Radio Challenge is an excellent opportunity for those in academic, government and industry research to interact with students whose designs represent the state-of-the-art in the software radio domain. We hope that this opportunity will be fully exploited in the future.

## VI. CONCLUSION

In this paper, we presented our smart radio entries for SDR Forum's Smart Radio Challenge for the years 2007 and 2008. Beginning with the 2007 challenge, we reviewed the challenge scenario and problem statement. An introduction of our smart radio system and a discussion of how the challenge problem shaped our design process followed. The system itself was presented first in overview and then block by block. Discussion of the challenge year finished with a review of the risks, obstacles, lessons and overall contributions. After repeating the discussion focusing on the 2008 challenge,

we concluded with some additional thoughts about the Smart Radio Challenge itself.

## REFERENCES

- [1] S. Midkiff and C. Bostian, "Rapidly-deployable broadband wireless networks for disaster and emergency response," in *The First IEEE Workshop on Disaster Recovery Networks (DIREN 02)*, 2002.
- [2] C. Bostian *et al.*, "Testbed for high-speed end-to-end communications in support of comprehensive emergency management," in *Proceedings of the 2004 annual national conference on Digital government research*. Digital Government Society of North America, 2004, pp. 1–2.
- [3] C. Rieser *et al.*, "Cognitive radio testbed: further details and testing of a distributed genetic algorithm based cognitive engine for programmable radios," in *Military Communications Conference, 2004. MILCOM 2004. IEEE*, vol. 3, Oct.-3 Nov. 2004, pp. 1437–1443 Vol. 3.
- [4] C. W. Bostian, "A prototype public safety cognitive radio for universal interoperability," Center for Wireless Telecommunications, Wireless at Virginia Tech, 2006.
- [5] J. Bernthal *et al.*, "Trends and precedents favoring a regulatory embrace of smart radio technologies," in *New Frontiers in Dynamic Spectrum Access Networks, 2007. DySPAN 2007. 2nd IEEE International Symposium on*, April 2007, pp. 633–648.
- [6] F. Ge *et al.*, "Cognitive radio: From spectrum sharing to adaptive learning and reconfiguration," in *Aerospace Conference, 2008 IEEE*, March 2008, pp. 1–10.
- [7] J. Krikke and B. Alfonsi, "In the news," *IEEE Intelligent Systems*, vol. 21, no. 2, pp. 4–7, 2006.
- [8] R. Rubenstein, "Radios get smart," *Spectrum, IEEE*, vol. 44, no. 2, pp. 46–50, Feb. 2007.
- [9] M. D. Silvius *et al.*, "Smart radio: spectrum access for first responders," in *Wireless Sensing and Processing III*, S. A. Dianat and M. D. Zoltowski, Eds., vol. 6980, no. 1. SPIE, 2008, p. 698008. [Online]. Available: <http://link.aip.org/link/?PSI/6980/698008/1>
- [10] M. D. Silvius, A. B. MacKenzie, and C. W. Bostian, "A survey of neighbor discovery protocols in modern wireless systems and their application to the design an implementation of 'rendezvous' algorithms in smart and cognitive radios," in *Virginia Tech Symposium on Wireless Personal Communications*, Blacksburg, VA, 2007.
- [11] SDR Forum. (2009) Sdr forum's smart radio challenge. [Online]. Available: <http://www.radiochallenge.org/>
- [12] —. (2007) 07 sample problems. [Online]. Available: <http://www.radiochallenge.org/07SampleProblems.html>
- [13] B. Le *et al.*, "A public safety cognitive radio node," in *Software Defined Radio Forum Technical Conference*, 2007.
- [14] M. D. Silvius, "Sdr forum's smart radio" challenge 2007 final report," 2007.
- [15] (2009) Gnu radio. [Online]. Available: <http://gnuradio.org/trac>
- [16] (2009) Ettus research llc. [Online]. Available: <http://www.ettus.com/>
- [17] Lyrtech. (2009) Small form factor sdr development platforms. [Online]. Available: <http://www.lyrtech.com/index.php?act=view&pv=SFF%20SDR%20development%20platforms>
- [18] R. Rangnekar *et al.*, "A remote control and service access scheme for a vehicular public safety cognitive radio," in *2009 IEEE Vehicular Technology Conference*, 2009, To Appear.
- [19] SDR Forum. (2008) 08 sample problems. [Online]. Available: <http://www.radiochallenge.org/08SampleProblems.html>
- [20] M. D. Silvius *et al.*, "Cwt smart radio system for smart radio challenge 2008," in *2008 Software Defined Radio Technical Conference, Poster Session*, 2008.
- [21] Q. Chen and C. W. Bostian, "Cognitive gateway design to promote universal interoperability," in *2008 Software Defined Radio Technical Conference*, 2008.

- [22] Q. Chen, Y. Wang, and C. Bostian, "Universal classifier synchronizer demodulator," in *Performance, Computing and Communications Conference, 2008. IPCCC 2008. IEEE International*, Dec. 2008, pp. 366–371.
- [23] K. Cheok, N. Loh, and H. Hu, "Self-training cognitive preview control for autonomous vehicle path navigation," in *Decision and Control, 1988., Proceedings of the 27th IEEE Conference on*, Dec 1988, pp. 2286–2291 vol.3.
- [24] F. Ge *et al.*, "A cooperative sensing based spectrum broker for dynamic spectrum access," in *Military Communications Conference. MILCOM 2009. IEEE*, 2009, submitted.
- [25] Y. Wang *et al.*, "Dynamic cellular cognitive system," November 2007.

## VII. ACKNOWLEDGEMENT

This work used technology and research results from projects supported by the National Institute of Justice, Office of Justice Programs, and U.S. Department of Justice under Award No. 2005-IJ-CKK017 and by the National Science Foundation under grant CNS-0519959. In addition, we would like to thank the Air Force Institute of Technology for its continued support of this work. The opinions, findings, and conclusions or recommendations expressed are those of the authors and do not necessarily reflect the views of the Department of Justice, the National Science Foundation, Department of Defense or US Air Force.