Design, Implementation and Characterization of Practical Distributed Cognitive Radio Networks

Ahmed Khattab, *Member, IEEE,* Dmitri Perkins, *Senior Member, IEEE,* and Magdy A. Bayoumi, *Fellow, IEEE*

***Abstract*—Opportunistic Spectrum Access (OSA) in distributed cognitive radio networks (CRNs) has been well studied in the literature from a theoretical perspective. However, such theoretically-optimized distributed OSA approaches are challenged by several practical implementation issues. In this paper, we design a custom cross-layer framework that enables the: *(i)* clean-slate implementation of a wide variety of OSA mechanisms; *(ii)* experimental evaluation of the individual practical OSA components of the Rate-Adaptive Probabilistic (RAP) framework; and *(iii)* detailed comparison of the performance of such a practical OSA approach against theoretical OSA approaches developed for fully-capable CRNs. Our evaluation reveals the multi-fold goodput improvement and remarkable fairness characteristics of the practical RAP OSA approach compared to the OSA approaches that overlook the OSA and CR practical limitations. However, the superior performance of practical OSA comes at the expense of more outages to the primary licensed networks but within the permissible bounds. Another key finding is that the wide family of existing theoretically-optimized OSA protocols can benefit from the gains available to the individual components of the practical RAP approach, namely, the random spectrum sensing and the probabilistic non-greedy access.**

***Index Terms*—Cognitive radio; radio spectrum management; ad hoc networks; design for experiments; performance evaluation.**

I. Introduction

OPPORTUNISTIC Spectrum Access (OSA) has recently been considered as a paradigm shift that will shape the future of wireless networks. OSA relies on cognitive radios (CRs) to exploit the temporally unutilized spectrum bands that are typically licensed to certain wireless services. The problem of optimizing the performance of a cognitive radio network (CRN) while providing grantees on the performance of the collocated primary licensed networks (PRNs) has received significant research interest from a theoretical perspective as discussed in [1], [2], and the references therein. Meanwhile, the implementation issues of OSA and CRNs have received less attention especially in distributed networks that lack centralized entities which involvement reduces the CRN complexity. Directly applying the theoretically-driven OSA approaches is challenged by several practical issues. Two crucial challenges that face distributed CRNs are: *(i)* the limited hardware capabilities given the stringent processing and speed requirements of CRNs (e.g., a low-cost multi-GHz transceiver does not exist yet), and *(ii)* the lack of a practical mechanism to measure the actual interference at primary receivers, and hence, relying only on the local measurement of the activities of the primary senders to make spectrum access decisions. Consequently, such networks are susceptible to making wrong spectrum access decisions [1]–[3].

In our prior work [3], we have presented the Rate-Adaptive Probabilistic (RAP) OSA framework that takes such challenges into account while formulating the OSA problem. RAP does not rely on explicit inter-flow coordination to avoid the associated overhead and the other challenges resulting from the use of a common control channel for global network coordination. However, RAP still uses a control channel to coordinate the spectrum decisions between a cognitive sender and its receiver. Using analysis and packet-level simulations, RAP OSA was shown to exhibit remarkable goodput and fairness characteristics compared to the hypothetically-optimal approaches which assume currently unavailable transceivers and explicit inter-flow coordination.

In this paper, our focus is on the less well-studied issue of implementing distributed OSA techniques given practical radio transceiver technologies and to characterize the performance of different OSA approaches in real systems. Our aim is not only to demonstrate the superior performance of the RAP approach in a real system but also to thoroughly investigate the role of the individual RAP components in the overall performance and their applicability to other existing OSA approaches. We use the Wireless open-Access Research Platform (WARP) [4] for our empirical study. WARP is well recognized by both the academic and industrial research communities for clean-slate prototyping. The contributions of the paper are:

- The design of a clean-slate OSA implementation framework using the WARP Field Programmable Gate Array (FPGA)-based platform that resembles practical wireless transceivers. Our design is modular to encapsulate the basic functions common to different OSA approaches such as spectrum sensing, common control channel, spectrum coordination packet handshake, and multi-rate multi-power packet transmission. Our OSA realization framework operates in realtime at bandwidths and time scales comparable to modern wireless devices. This contrast with existing Software-Defined Radio (SDR)-

Manuscript received January 22, 2013; revised May 22 and July 25, 2013. The editor coordinating the review of this paper and approving it for publication was Y. J. (A.) Zhang.

A. Khattab is with the Department of Electronics and Electrical Communications Engineering, Cairo University, Egypt (e-mail: akhattab@ieee.org).

D. Perkins and M. A. Bayoumi are with the Center for Advanced Computer Studies (CACS), University of Louisiana at Lafayette, LA, USA (e-mail: {perkins, mab}@cacs.louisiana.edu).

The work was supported in part by the National Science Foundation (NSF) under NSF Career Grant No. 0448055, the U.S. Department of Energy (DOE) under Award Number DE-FG02-04ER46136, and by the State of Louisiana, Louisiana Board of Regents under Contract Numbers DOE/LEQSF(2004-07)-ULL and LEQSF(2003-06)-RD-A-35.

Digital Object Identifier 10.1109/TCOMM.2013.090513.130060

based OSA implementations [5]–[7] which have limited bandwidth operation due to platform interfacing bottlenecks (e.g., the maximum transmission rates of the major SDR OSA implementations [7] and [5], [6] are 1 Mbps and 1.5 Mbps, respectively). Furthermore, our clean-slate implementation is not based on the IEEE 802.11 medium access control (MAC) as the case with the WARP-based OSA implementation presented in [8] or existing OSA platforms that are based on commodity hardware such as [9], [10]. A thorough overview of the related OSA implementations will be presented in Section II.

- The complete implementation details and performance evaluation of the practical RAP OSA approach using our implementation framework. We also implement a set of OSA protocols that serve as benchmarks. A given OSA protocol can easily implement its spectrum management approach through our framework's flexible function interface designed to facilitate the interaction with the basic functions. Our modular OSA design allows us to separate and reuse the key components of the RAP approach, namely, the random sensing component and the non-greedy probabilistic access component. Hence, we can easily integrate such modules in the implementation of other OSA approaches to evaluate the potential performance gains if such modules are used instead of their existing counterparts developed without regard to the practical limitations of existing CR technologies.
- An extensive set of experiments detailing the performance of the RAP OSA framework.[1] Our experiments isolate and explain the contribution of the individual practical components of the framework. We also evaluate the performance of a representative candidate of the wide set of existing OSA approaches that are solely based on theoretical assumptions. Our results demonstrate the substantial performance gains that are available to such a wide range of existing OSA approaches when using the individual components of our practical OSA framework.

The remainder of the paper is organized as follows. We motivate our work by reviewing the related OSA implementation literature in Section II. Then, we describe our hardware implementation framework in Section III. We briefly present the RAP OSA approach as well as the other OSA protocols used in our performance evaluation in Section IV and present the results of an extensive set of experiments in Section V. We conclude the paper in Section VI.

II. Related Work

A. SDR-Based OSA Implementations

In his seminal work [13], Mitola defined the Software-Defined Radio (SDR) as the ideal environment for implementing CRs due to its seamless flexibility and controllability. An SDR platform relies on a low-cost Universal Software Radio Peripheral (USRP) [14] interfaced to a general purpose computer running the software development environment: the GNU radio [15]. Several OSA implementation frameworks have been developed to provide the software libraries and environments enabling the fast composition of OSA protocols. Examples include the Cognitive Radio Open Source System (CROSS) [5] and its hardware platform: the COgnitive Radio NEtwork Testbed (CORNET) [6], the Papyrus software platform [7] and the Sora platform [16], as well as the platforms presented in [17], [18]. However, such SDR platforms have a major shortcoming in their transmission speeds (only few Mbps as in [5]–[7]) due to the slow interface between the USRP and the general purpose processor as comprehensively discussed in [19]. Recent cognitive SDR-based platforms such as [20] and [16] are aiming at providing the ability to dynamically add hardware-based acceleration for addressing the SDR latency and bandwidth limitations. Alternatively, [21] propose a novel architecture that replaces the SDR frontend with an RF Application Specific Integrated Circuit (ASIC) to improve the CR speed. However, the performance of all such platforms is way far from the performance of real-life systems.

B. Commodity Hardware-Based OSA Implementations

Several attempts have been made to use commodity IEEE 802.11 hardware to implement CR terminals [22]–[25]. While such an implementation approach provides a low-cost solution, it has a limited room for reconfigurability and customization since such an IEEE 802.11 hardware is restricted in providing accessability to the underlying physical layer parameters. Alternatively, [9] allows the node to choose the approximate "best-fit" MAC from a limited number of predefined stand-alone MAC protocols that are implemented using an IEEE 802.11 hardware. Furthermore, [10] proposed a modular approach to break down the different MAC approaches into a few modular designs. However, all such platforms lack the flexibility of SDR-based platforms and are incapable of satisfying the requirements of fully-capable CRNs.

C. FPGA-Based OSA Implementations

In contrast, Field Programmable Gate Array (FPGA)-based platforms provide both the flexibility of SDR-based platforms and the practical performance of hardware-based platforms. However, such remarkable performance of FPGA-based platforms comes at the expense of increased complexity and cost. An example FPGA-based prototype is the network-centric cognitive radio (WiNC2R) that is equipped with a tri-band radio frontend operating at the 700 MHz, 2.4 GHz and 5 GHz bands [26]. Likewise, DARPA is currently developing a hand-held FPGA-based CR terminal covering the frequency range between 900 MHz and 6 GHz as a part of the Wireless Networks after Next (WNaN) program [27]. The WNaN program aims at reducing the transceiver cost to the point where a sophisticated, multi-transceiver CR can achieve a cost point below that of conventional technology [28]. Unlike the above two FPGA platforms specifically designed for cognitive radio applications, general purpose wireless FPGA platforms can be used for the implementation of OSA protocols. For example, the authors of [8] present an OSA implementation framework using the general purpose Wireless open-Access Research Platform (WARP) [4]. However, that framework is a derivative of the WARP IEEE 802.11-like MAC. In contrast,

[1] While [11] and [12] presented a preliminary subset of the experimental results and a tutorial-level overview of the implementation framework, respectively, here we present comprehensive implementation details and evaluations.

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

we next present the first implementation and performance evaluation of clean-slate OSA approaches using WARP.

III. Opportunistic Spectrum Access Implementation Framework

This section describes the first clean-slate OSA framework and the design steps in implementing different OSA protocols.

A. Wireless Open-Access Research Platform (WARP)

WARP is an FPGA-based hardware platform with an open-source repository of building blocks and reference designs [4]. WARP is ideal for clean-slate medium access prototyping through a flexible interface between the physical and medium access layers. WARP implements an OFDM transceiver on the FPGA fabric. We use the WARP OFDM physical layer implementation available in the WARP reference design unaltered. We build a new MAC framework for WARP that is generic for implementing OSA protocols. Our implementation approach is clean-slate and does not adopt the IEEE 802.11-like WARP MAC that comes with the WARP reference design that was used in [8] in the first WARP-based OSA implementation. Our OSA implementation framework is written in C-language, compiled and downloaded to one of the PowerPC cores of a WARP board where it directly interacts with the physical layer.

B. OSA Control Mechanisms

Our implementation instruments the basic four functionalities common to different OSA schemes in a clean-slate WARP MAC layer implementation. These common functions are: *(i)* spectrum sensing, *(ii)* common control channel, *(iii)* spectrum coordination packet handshake, and *(iv)* configurable multi-rate multi-power packet transmission. Our OSA MAC implementation runs on top of the WARP OFDM physical layer implementation included in the WARP reference design version 14 to directly read the RSSI values and configure the transceiver parameters.

- **Spectrum Sensing:** The function `SpectrumSelect(PROTOCOL)` measures the cumulative interference of a given spectrum band and determines whether it is below the power mask specified by the corresponding PRN or not. This is realized by monitoring the received signal strength indicator (RSSI) averaged over a certain time window. By comparing the time-averaged RSSI with the spectrum power mask, an OSA protocol can determine whether this band is clear (RSSI $<$ Power Mask) or not (RSSI $\geq$ Power Mask).
- **Common Control Channel:** Distributed OSA protocols require a means by which a cognitive sender coordinates its spectrum decisions with its intended receiver. A common control channel (CCC) is generally used for this purpose. Both the senders and the receivers are continuously listening to this channel if not involved in an active data exchange. We define channel 14 of the 2.4 GHz ISM band as the CCC. Channel 14 of the 2.4 GHz band is not available for commercial purposes in the United States and can only be used for academic research. Using such a channel guarantees a robust CCC.

TABLE I
OSA Implementation Framework Parameters.

Parameter	Value
Transceiver TX/RX Turnaround Delay	23 μsec
Channel Switching Time	10 μsec
Single Channel Sensing Time	30 μsec
Timeout Period	50 μsec
Payload Packet Length	1450 Bytes
Control Packet Length	24 Bytes
TX Power (BPSK, QPSK, 16QAM)	(12,15,18) dB
TX MAC Rate (BPSK, QPSK, 16QAM)	(4.1, 8.4, 10.4) Mbps

- **Spectrum Coordination Packet Handshake:** We create the control packets to be exchanged over the CCC for cognitive sender-receiver coordination. These control packets do not include any payload bytes and only include the sender and the intended receiver addresses in addition to other protocol-dependent control information such as the selected spectrum, the measured RSSI, the modulation rate, etc. RAP and the other tested OSA protocols only need a two-way control-message handshake in which the sender informs its receiver with its spectrum selections via one packet and the receiver confirms or denies such selections with another packet. The control packet handshake, depicted by Code 1, is transmitted using the default rate realized via WARP QPSK modulation scheme.
- **Configurable Packet Transmission:** Finally, we implement a packet transmission function `set_pkt_rate_and_power(PROTOCOL)` which configures the modulation rate and power parameters on a packet-per-packet basis. We allow an OSA protocol to configure the transmission channel, the modulation rate and power. A data packet can use one out of three WARP modulation schemes: BPSK, QPSK, and 16 QAM with respective transmission powers of 12 dBm, 15 dBm, and 18 dBm. Table I summarizes the main parameters of our implementation.

C. OSA Implementation Framework Overview

The above functionalities are common to different OSA protocols – despite the fact that each protocol adopts a different mechanism to manage the spectral opportunities. We encapsulate these basic functionalities and create a flexible function interface to facilitate the interaction with them to realize a generic framework for implementing OSA protocols. The flow chart shown in Figure 1 outlines the state machine of the proposed framework. The CR is set to be continually monitoring the CCC at the default operation mode. Two events can interrupt this default operation mode: the reception of a packet to be transmitted from higher layers through the function `dataFromNetworkLayer_callback()` and the correct reception of spectrum coordination packet destined to the node through the function `phyRx_goodHeader_callback()`. In the former case, the transceiver enters the transmit path depicted by the left hand side part of Figure 1 which invokes the `SpectrumSelect(PROTOCOL)` function which decides which spectrum to use, exchanges the spectrum coordination messages, configures the transmission parameters, and transmits the data packet accordingly. In the latter case, the transceiver

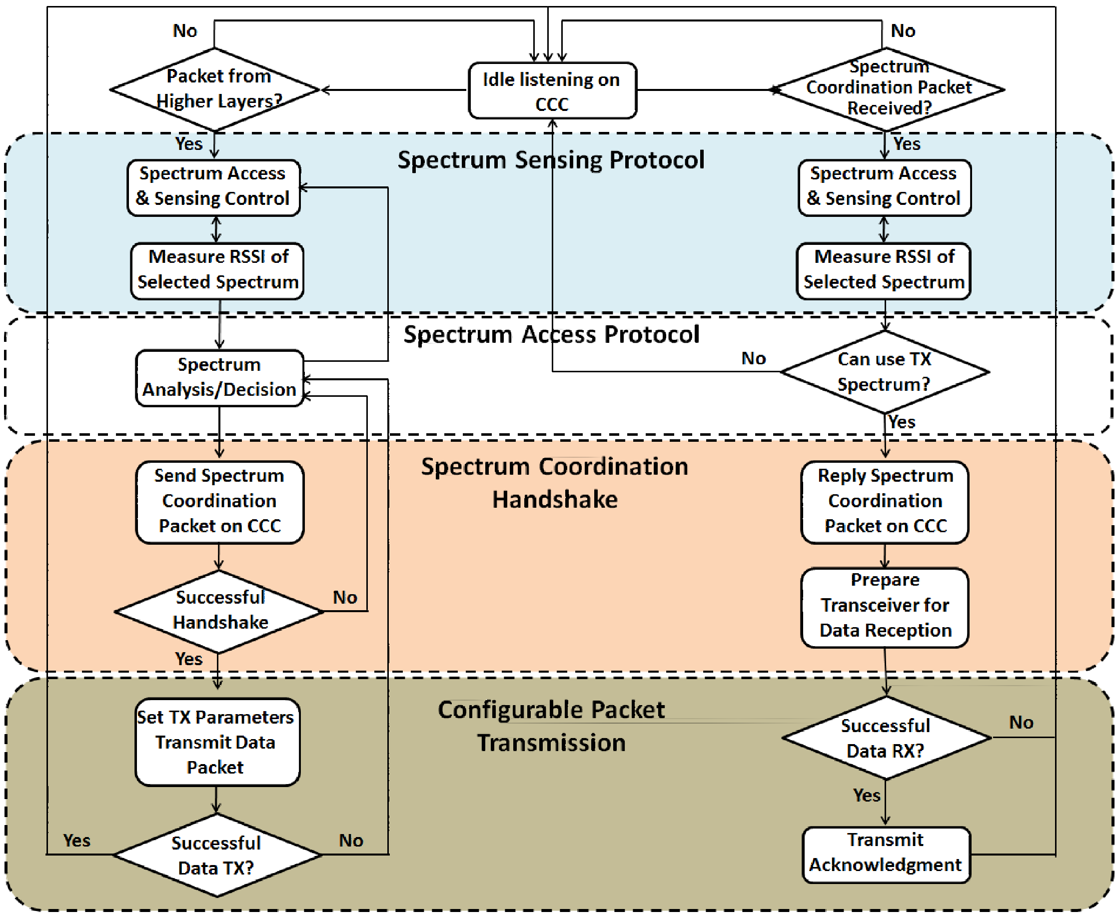

Fig. 1. Our generic clean-slate OSA implementation framework.

enters the receive path depicted by the right hand side part of Figure 1 which may run any spectrum access or sensing functions (based on the OSA protocol) before replying to the received spectrum coordination message and receiving the actual data packet. Code 1 and Code 2 depicts the pseudocode of the handler invoked in response to the reception of different packet types within `phyRx_goodHeader_callback()` and the timeout timer handler, respectively.

Code 1. Received Packet Handler Pseudocode

```
if(packet->header.destAddr == myID){
    switch(packet->header.pktType){
      case SPECTREQ:
        chan = packet->header.channelID;
        if(RX_SPECT_SENS_REQUIRED) {
          warpphy_setChannel(chan);
          usleep(ChannelSwitchTime);
          SpectStatus = warpmac_carrierSense();
        }
        warpphy_setChannel(CNTRLCHANNEL);
        usleep(ChannelSwitchTime);
        set_pkt_rate_and_power(default);
        send_spectrpl(SpectStatus);
        warpphy_setChannel(chan);
        usleep(ChannelSwitchTime);
        warpmac_setTimer(TIMEOUT_TIMER);
        break;
      case SPECTRPL:
        if(warpmac_inTimeout()){
          warpmac_clearTimer(TIMEOUT_TIMER);
          chan = packet->header.channelID;
          set_pkt_rate_and_power(PROTOCOL);
          warpphy_setChannel(chan);
          usleep(ChannelSwitchTime);
          send_data();
          warpmac_setTimer(TIMEOUT_TIMER);
          warpmac_decrementRemainingReSend();
        }
        break;
      case DATAPACKET:
        warpmac_clearTimer(TIMEOUT_TIMER);
        state = warpmac_finishPhyRecv();
        if(state & PHYRXSTATUS_GOOD)
          send_ack();
        warpphy_setChannel(CNTRLCHANNEL);
        usleep(ChannelSwitchTime);
        break;
      case ACKPACKET:
        if(warpmac_inTimeout()){
          warpmac_clearTimer(TIMEOUT_TIMER);
          txMacframe.remainingTx = 0;
          set_favorite_spectrum();
          warpmac_enableDataFromNetwork();
        }
        break;
    }
}
```

Code 2. Timeout Timer Handler Pseudocode

```
if (txMacframe.header.length < THRESHOLD) {
  warpphy_setChannel(CNTRLCHANNEL);
  usleep(ChannelSwitchTime);
}
else {
  warpphy_setChannel(SpectrumSelect(PROTOCOL));
  usleep(ChannelSwitchTime);
}
if (TIMEOUT_TIMER_RX_MODE){
  warpphy_setChannel(CNTRLCHANNEL);
  usleep(ChannelSwitchTime);
  TIMEOUT_TIMER_RX_MODE = 0;
  warpmac_enableDataFromNetwork();
  break;
}
else {
  warpmac_setTimer(TX_TIMER);
}
```

Note that the common control channel, spectrum sensing, spectrum coordination handshake and configurable packet transmission functions in the framework are independent of the particular OSA protocol. Meanwhile, a particular OSA controls such functions. For instance, the OSA protocol defines which spectrum bands to be sensed by the sensing module, which spectrum to use, and what are the appropriate transmission parameters in the generic framework shown in Figure 1. Our modular implementation framework provides a flexible interface to the basic functions that allows for the implementation of a wide variety of OSA protocols. Next, we use this generic framework to implement the set of OSA protocols used for our empirical study.

IV. Towards Practical Distributed Opportunistic Spectrum Access

In this section, we overview the main components and implementation details of the practical RAP OSA approach.[2] Then, we discuss the other implemented OSA approaches we use as benchmarks for our performance evaluation study.

A. Practical Challenges of Distributed OSA

1) CR Hardware Limited Capabilities: OSA requires the CR to be able to identify the spectral opportunities and take decisions regarding the most appropriate band to use. This poses stringent requirements on the transceiver hardware. For instance, a CR transceiver should be able to sense the radio activities over a wide range of the spectrum (e.g., multiple gigahertz) and process the acquired measurements to come up with access decisions in very short time intervals. While significant efforts are currently being made to provide such high performance transceivers at low cost, existing technologies still cannot allow exploiting cognitive radio networking at its full potential as discussed in Section II in details. Existing CR hardware is either *(1)* having orders of magnitude lower communication speeds compared to real-life systems (SDR-based CRs), *(2)* not fully reconfigurable (CRs based on 802.11 hardware), or *(3)* expensive (FPGA-based CRs). The gap between existing hardware technologies and the performance requirements of CR motivated FCC to drop the spectrum sensing function of the white space CR devices in 2010, while urging the development of highly capable CR hardware [29].

2) Distributed OSA Coordination Problem: OSA necessitates that the CR users do not disturb the transmissions of the primary users. Furthermore, CR users within a CRN should coordinate their transmissions in order to share the available spectral opportunities. Several opportunistic spectrum management schemes have been proposed in the literature aiming at optimizing the CRN goodput while not degrading the performance of the primary networks [30]–[39]. However, such schemes do not take into account the practical limitations of CRNs. More Specifically, the CRN users are not able to measure the interference at the primary receivers, and instead, take their spectrum access decisions based on their local measurements of the transmissions of the primary senders. Even with more accurate channel quality assessment techniques, such as that presented in [40] for IEEE 802.11 channels, local assessment of the channel quality does not reflect the surrounding receivers. Such inaccurate sensing, combined with the speed and other hardware limitations of cognitive transceivers, does not guarantee the optimality of existing schemes in maximizing the CRN goodput or satisfying the PRN constraints. Furthermore, such schemes adopt greedy spectrum access mechanisms in order to maximize the utilization of any available spectral opportunities, and hence, require an explicit control mechanism for coordinating the spectrum decisions of different competing cognitive transmissions. Such greedy strategies lead to unfairness in the CRN goodput distribution and cause the control mechanism to be a bottleneck and a single point of failure of the system which is strongly undesirable [1], [2].

B. RAP Random Sensing with Probabilistic Access OSA

We presented the Rate-Adaptive Probabilistic (RAP) spectrum access approach for low-complexity and practical cognitive radio networks in [3]. The RAP approach is developed to counter the CR limitations and the consequent unavoidable inaccuracy in spectrum sensing and to avoid the overhead of explicit inter-flow coordination. Furthermore, RAP addresses the challenge wherein the PRN performance should not be hurt in the absence of spectrum sensing as in FCC white space systems [29]. The RAP approach adopts coordinated random spectrum selection combined with a rate-adaptive probabilistic transmission policy to achieve the above goals.

1) Coordinated Random Sensing: The coordinated random spectrum selection component has the CR sender randomly selecting a spectrum to probe for an upcoming transmission due to its inability to accurately assess the impact of its transmission on the primary receivers. This relaxes the requirements of the sensing module of a cognitive radio (since it does not require a wide-band front-end). Furthermore, the RAP approach has both the flow's endpoints participate in the decision of whether or not to use the randomly-selected spectrum by having both the CR sender and receiver measuring the interference on the sender-selected spectrum. Unlike traditional listen-before-talk OSA and MAC protocols, the

[2] A detailed description the practical RAP OSA approach and its analytical optimization is available in [3].

RAP transmission decision takes into account the spectrum conditions at the receiver side which typically differ from the spectrum conditions at the sender side in ad-doc networks wherein nodes are exposed to different interference conditions. The CR sender-receiver pair will only utilize the randomly selected spectrum if the receiver measurement indicates a clear spectral opportunity regardless of the sender measurement. The sender measurement is used in the decision of the transmission scheme as explained next.

2) Rate-Adaptive Probabilistic Access: The rate-adaptive probabilistic component has the sender probabilistically deciding whether or not to transmit and at which rate and power based on the two interference measurements at both flow's endpoints and the fact that these measurements do not actually reflect the interference at nearby primary receivers. The RAP access policy is non-greedy in the sense that the highest rate and power applicable to a given scenario are not used with a unity probability. Instead, lower rates and powers are to be probabilistically used to *(i)* counter the unavoidable inaccuracy in spectrum sensing due to hidden and exposed primary nodes (since spectrum sensing techniques only measure the transmission activities of the primary senders), and *(ii)* prevent a single cognitive sender-receiver pair from monopolizing a spectral opportunity to allow multiple flows to share such an opportunity, and hence, alleviates the need for explicit inter-flow coordination.

This component is implemented as follows. The RAP sender will use either the highest possible rate/power with probability p or a variable lower rate/power with probability $1-p$, where p is the RAP probability of transmission at the highest possible rate/power when the interference measurements at both flow's endpoints are below the power mask of the spectrum. Alternatively, the RAP sender will either randomly select a new spectrum to use with probability $1 - q$ or use the minimum rate/power with probability q , where q is the RAP probability of transmission when the sender's interference measurement exceeds the power mask defined by the primary owner of the spectrum while the receiver interference measurements is below the power mask. Further details of the RAP framework and protocol details are available in [3].

We analytically formulated the constrained CRN performance optimization problem as a mixed-integer non-linear program to derive the optimal values of the RAP OSA approach. For the rest of the paper, we refer to the RAP approach in terms of its two main components as random sensing with probabilistic access. Our goal is to identify the performance gain of each component and how they can be jointly or individually used to improve the performance of existing OSA approaches that are theoretically optimized without regard to the practical limitation of OSA in distributed CRNs.

C. RAP Protocol Implementation

RAP-MAC is the medium access control protocol implementation of the RAP OSA approach. We implement the RAP-MAC using the implementation framework presented in Section III. We write the RAP-MAC sender and receiver state machines, depicted by the flow chart in Figure 2, in C-language and compile them in a single bit file to reside on one of the WARP FPGA PowerPC cores. A demonstration of the RAP-MAC implementation in action is available online [41]. The RAP-MAC implementation simply integrates the two main RAP components with the generic functions of our general implementation framework as shown in Figure 2. A RAP node is continuously listening to the common control channel. If the application layer of the node has a packet to be transmitted, the RAP node switches to the sender state machine depicted by the left hand side part of Figure 2. Alternatively, if the RAP node receives a SR packet destined to it while being in the idle listening state on the CCC, the node moves to the receiver state machine depicted by the right hand side part of Figure 2. A successful RAP sender-receiver data packet exchange is described as follows.

- Upon receiving a packet to transmit, a RAP sender randomly picks a spectrum band – if a favorite spectrum that recently carried out a successful transmission does not exist – to use, measures the interference of this band, and transmits a spectrum request (SR) packet containing the spectrum identity and the RSSI measurement.
- When the targeted receiver receives the SR packet, it also measures the interference over the sender-selected spectrum. If the receiver can successfully receive packets over the selected spectrum, it replies with a spectrum grant (SG) packet that includes the maximum rate it can sustain given its RSSI measurement as well as an indicator of whether or not this spectral opportunity is clear. Otherwise, the RAP receiver will return to the CCC idle listening state if it cannot use the sender-selected spectrum given its interference assessment.
- If the sender receives the SG packet within the timeout period (50 microseconds), it uses its own assessment of the chosen spectrum alongside the receiver-side information conveyed through the SG packet to decide whether or not to use this spectrum and the appropriate transmission power/rate according to the RAP probabilistic access component discussed above. If the timeout timer expires before the reception of the SG packet, the RAP sender randomly selects a new spectrum if the packet is not to be discarded yet.
- If the RAP receiver – which has tuned itself to the sender-selected spectrum – correctly receives the data packet, it acknowledges the packet reception. Then it returns to the CCC idle listening state.
- Upon receiving the acknowledgement within the timeout period, the RAP sender will decide whether or not to declare this spectrum as its favorite for upcoming transmissions (in the case a clear spectral opportunity) before returning to the CCC idle listening state.

D. Benchmark Protocols Implementations

Our objective is not only to demonstrate the advantages of the RAP OSA approach but also to study how much gain is attributed to its different practical mechanisms. We also show how traditional OSA approaches can benefit from the individual RAP components. Consequently, we implement the following suite of OSA protocols for our experimental study.

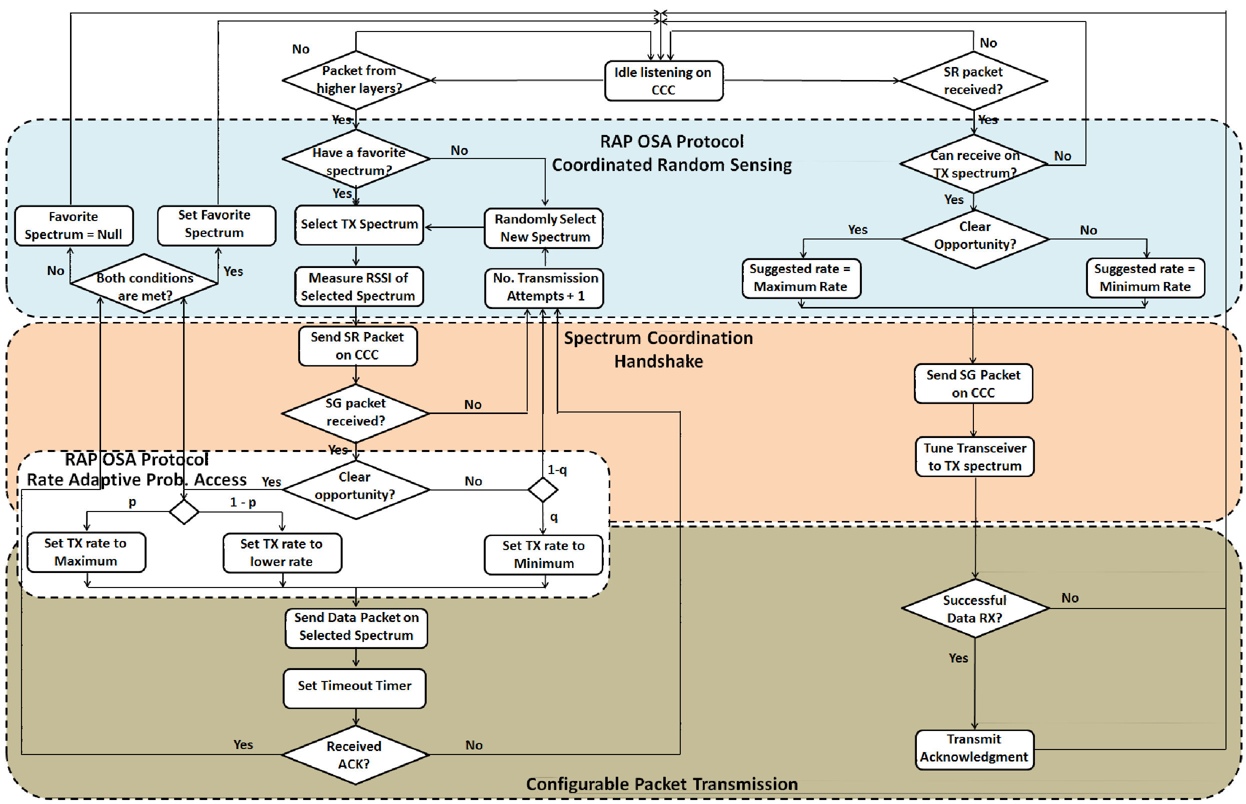

Fig. 2. Detailed RAP implementation using the proposed OSA implementation framework.

1) Sequential Sensing with Greedy Access: This implementation reflects the wide range of existing OSA protocols (e.g., [36]–[39]) that are theoretically optimized without regard to the practical CRN limitations. In such schemes, a cognitive radio node senses all of the available spectrum bands before deciding which band to use. Unlike the RAP approach, such schemes adopt deterministic and greedy (i.e., winner-takes-all) access mechanisms in which a sender only transmits if there exists a spectrum which its measured RSSI is below the power mask. Furthermore, such senders transmit using the highest possible power and rate for all the time. We use a modified version of a candidate protocol of this family of protocols that was presented in [36] for our implementation. The spectrum access mechanism of such protocols is based on traditional carrier sensing that uses a two-way message exchange over the CCC to insure a single secondary user transmission per contention area. However, we do not implement the ability of sensing or transmitting over multiple bands simultaneously as in the original protocol presented in [36] to adapt to the limited capabilities of conventional transceivers and for the sake of fairness in comparison. Recall that the WARP transceiver can be tuned to only one frequency channel at a time as the case with other contemporary single-radio transceivers. Therefore, we implement a sequential spectrum sensing mechanism in which a cognitive node goes over the channels of interest and reports back the RSSI of individual channels.

2) Sequential Sensing with Probabilistic Access: The second protocol that we use for comparison is a derivative of the above implementation which still depends on sequentially scanning all of the available spectrum bands before deciding the best spectrum to use based on traditional carrier sensing. However, this protocol adopts a probabilistic and non-greedy spectrum access approach similar to that developed for the RAP OSA approach instead of using deterministic and greedy spectrum access. Such a protocol helps identifying how much gain can be achieved by using a probabilistic access mechanism if adopted by the wide range of existing protocols that rely on greedy access strategies. Furthermore, this protocol implementation allows us to assess how much gain of the RAP approach is due to random sensing since the sensing mechanism is the only difference between RAP OSA approach and this protocol implementation.

3) Random Sensing with Greedy Access: We also implement a variant of the RAP-MAC protocol in which RAP probabilistic access is replaced with traditional winner-takes-all access. We refer to this implementation as the random sensing with greedy access protocol. The greedy access mechanism of this protocol is the same one used by the sequential sensing with greedy access protocol. Hence, this protocol allows us to quantify the performance gain of randomized narrow-band sensing compared to sequential wide-band sensing. Furthermore, comparing the performance of this protocol against RAP-MAC illustrates the contribution of the RAP probabilistic access component in the overall RAP gain as will be demonstrated next by our experiments.

V. Empirical Performance Evaluation

In this section, we present an exhaustive set of experiments that does not only explain the superior performance of practical OSA mechanisms but also demonstrates that existing theoretical-driven OSA approaches can benefit from the gains of the individual practical mechanisms.

A. Experimental Methodology

Implementing a CRN environment poses significant design challenges. An OSA experiment requires the creation of multiple PRNs which spectral opportunities can be exploited by the CR users when the primary users are inactive. Thus, the experiments must provide controllable PRN flows. Furthermore, the experimental setup must keep track of every CRN transmission as well as every transmission and reception for all the PRNs in order to assess the CRN decision mechanism and the outage performance of the PRNs, respectively.

Primary Networks Implementation: For our experiments, we create two PRNs each composed of a single sender and a single receiver. In order to have full control over the performance of the PRNs and to not harm existing licensed networks, we configure the two PRNs to operate over non-overlapping channels in the unlicensed 2.4 GHz ISM band. More specifically, we configure the first PRN to use channel 1 of the 2.4 GHz and the second PRN to use channel 7 of the same band. We use laptops equipped with IEEE 802.11 Ubiquiti SRC 802.11g/b/a wireless PCMCIA cardbus configured in the 802.11b ad-hoc mode to create the PRNs.[3] We set the transmission power of the PRNs to 18 dBm as the WARP nodes and disable the auto-rate fall back feature to maintain the transmission rate at its maximum value (11 Mbps). We use *iperf* to generate a UDP flow from each primary sender and collect the UDP flow statistics at the corresponding receiver. We measure the backlog UDP capacity of the two PRNs in the absence of any CRN activities to be 6.03 Mbps and 6.15 Mbps, respectively. Note that due to the difference in the noise floor of the used channels, there are slight differences in their goodput and the outage performances – despite having similar trends – as experienced in our experiments.

Cognitive Radio Network Implementation: We create a cognitive radio node by connecting a laptop (with its wireless interface disabled) to a WARP board via the WARP Ethernet port. By downloading the appropriate bit file of any of the implemented opportunistic spectrum access protocols to a WARP PowerPC, the WARP board will act as the wireless air interface of the laptop that runs that particular OSA protocol. We create a fully backlogged CR transmission between two such CR nodes using *iperf*. The CR sender and receiver nodes are at equal distance of approximately 2 meters from the senders and receivers of the two collocated primary networks. Figure 3 depicts a layout of the experiment setup.

Performance Metrics: Our performance metrics are both the goodput of the CRN flow (defined as the amount of data

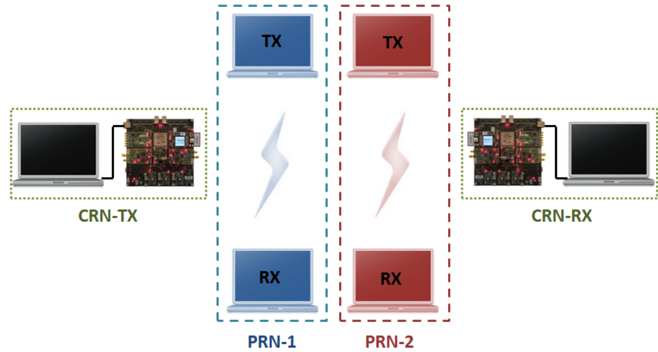

Fig. 3. Illustration of the experiment setup.

correctly received at the receiver) as well as the outage probability of both PRNs (defined as the percentage of the loss in the transmission rate due to the activity of the CRN). The reported results in the next subsection are the average of several runs each of one minute length. We run the experiments between midnight at the early hours of the morning to minimize the potential uncontrolled transmissions over the used channels.

B. Experimental Results

1) RAP Parameter Characterization: We start by characterizing the performance of the RAP approach in the worst-case scenario in which the PRNs are fully utilized. Our goal is to identify the values of the parameters of the RAP probabilistic access component. Note that the optimal parameter values analytically derived in [3] for arbitrary large-scale networks do not directly apply to our testbed setup due to the difference in the underlying system assumptions. We perform a two-dimensional sweep of the probability of transmission when the spectrum is clear, p, and unclear, q. As shown is Figure 4, the transmission probabilities p and q that achieve the highest CRN goodput for a targeted 5% maximum PRN outages are 0.4 and 0.4, respectively. We use these values for the rest of our experiments. Beyond these values of p and q, our measurements showed that the CR flow goodput decreases then increases again as q approaches unity with p having only a slight impact on the performance. This is because the probability of colliding with a primary transmission increases with the cognitive sender attempting more aggressively to exploit such unclear spectral opportunities. Such a negative impact of increasing the value of q does not only degrade the performance of the CR transmission but also negatively impact the outage of the PRNs as illustrated in Figure 4. Such significant PRN outages caused by the high q value will cause the IEEE 802.11 backoff window of both PRN flows to increase. Such large backoff windows create non-authentic spectral opportunities that are exploited by the CR flow to obtain a high goodput.

2) CRN Goodput Performance: Figure 5 illustrates the CRN flow goodput according to different OSA protocols. We vary the activity factor of both PRNs by varying the UDP flow rate such that the PRN activities go from idle to fully backlogged in 25% activity increments. As shown in Figure 5(a), RAP-MAC achieves the highest cognitive flow goodput among the four implemented OSA protocols for different

[3] Our use of IEEE 802.11 PRNs does not affect the generality of our results because the implemented OSA protocols do not compete with the PRNs for channel access. This is because the time needed to access a channel in the implemented OSA protocols is much larger than the DIFS period of the used IEEE 802.11 cards due to channel switching and other WARP delays.

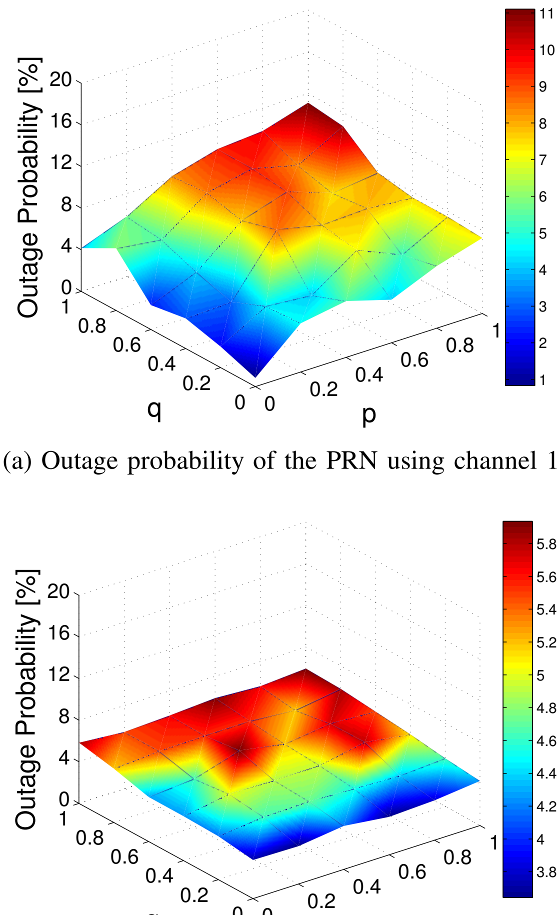

(a) Outage probability of the PRN using channel 1.

(b) Outage probability of the PRN using channel 7.

Fig. 4. Worst-case outage probability of the primary networks. For a worst-case outage of 5%, the optimal p and q values are 0.4 and 0.4, respectively.

PRN activities except when the PRNs are idle, in which case the goodput of the random sensing greedy access approach is slightly higher than the RAP-MAC goodput. Meanwhile, the sequential sensing with greedy access approach widely used for OSA results in the lowest goodput. The RAP-MAC goodput gain increases from 66% at low PRN activities to 95% at 50% PRN activity as shown in Figure 5(b). As the PRN activities increase, the RAP-MAC goodput becomes multiple folds of the goodput achieved by the benchmark protocol until the RAP-MAC goodput is 6.7 times the benchmark goodput when the PRNs are fully backlogged.

The superior RAP-MAC goodput performance is attributed to both its main components: the randomized sensing (which alleviates the overhead of scanning all frequencies before a given access by measuring the interference on only one randomly-selected frequency) and the non-greedy probabilistic access (which probabilistically explores the spectral opportunities rather than adopting the traditional winner-takes-all approach). We use the goodput achieved by the other two implementations (shown in Figure 5(a)) to perform pairwise comparisons to identify how much each component is contributing to the overall gain. Intuitively, adopting greedy access results in a slightly higher goodput when the PRNs are idle, regardless of the sensing mechanism. Hence, random sensing is the main contributor to the overall gain at low PRN activities as seen by comparing the protocols implementing random sensing against sequential sensing for both access mechanisms. More specifically, by comparing the solid blue curve (representing RAP-MAC) against the black dotted curve (representing its sequential sensing counterpart), and comparing the dashed light blue curve (representing random sensing with probabilistic access) against the dash-dotted red curve (representing the common greedy access based on sequential sensing). As the activities of the PRNs increase, the gain due to probabilistic access increases. For PRN activities of 50% and above, the contribution of random sensing is approximately 70% to 80% of the overall RAP-MAC gain while the contribution of the probabilistic access mechanism is around 20% to 30% depending on the PRN activities.

The relative gain of different OSA approaches with respect to the traditional approaches that use sequential sensing with greedy access depicted in Figure 5(b) roughly reflects the above percentages. More specifically, Figure 5(b) shows the gains available when the individual RAP components are used to improve the performance of traditional OSA. For example, adopting non-greedy access results in goodput gain of up to 56% as depicted by the dotted red curve. Furthermore, exploiting random sensing instead of sequentially searching for the best channel to use achieves 64% to 82% of the RAP gain, depending on the PRN activities, as illustrated by the light-blue dashed curve. This emphasizes that the random sensing component has a more significant performance gain. It is worth mentioning that while random-sensing-based OSA protocols randomly pick the used spectrum, they tend to have spectrum utilization patterns similar to OSA schemes that sense the entire spectrum as we have shown in [12].

3) PRN Outage Performance: Next, we evaluate the outage performance of the PRNs for different OSA protocol implementations. Two observations can be made regarding the PRN outages shown in Figure 6. First, probabilistic access schemes result in slightly higher PRN outages compared to their greedy access counterparts. However, probabilistic access has a weaker impact on the PRN outage when sequential sensing is used (as illustrated by the small gap between the dotted black and dash-dotted red outage curves). With the inaccuracies of random sensing, the impact of probabilistic access increases (as illustrated by the gap between the solid blue and dashed light-blue outage curves). Second, random sensing results in approximately 2.6 times the outages due to sequential sensing protocol irrespective of the access protocol. This is because sequential sensing protocols assess the interference levels on both channels before deciding the transmission action. On the other hand, random sensing protocols simply pick a channel at random for transmission. Note that despite resulting in higher PRN outages, random sensing protocols including RAP-MAC adhere to the targeted 5% maximum outage constraint. However, the significant multi-fold goodput gain of such protocol illustrated in Figure 5 outweighes the consequent excess primary outages. Furthermore, as the number of the primary networks increases, the sensing time required to assess the interference on all channels will increase. Hence, the RAP goodput gain is expected to further increase.

To conclude, the RAP approach improves the goodput

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

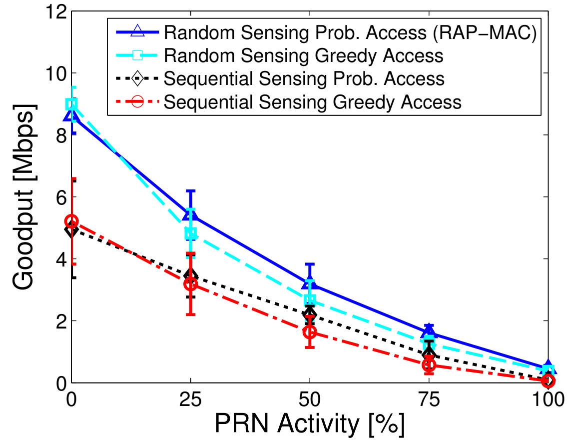

(a) CRN flow goodput.

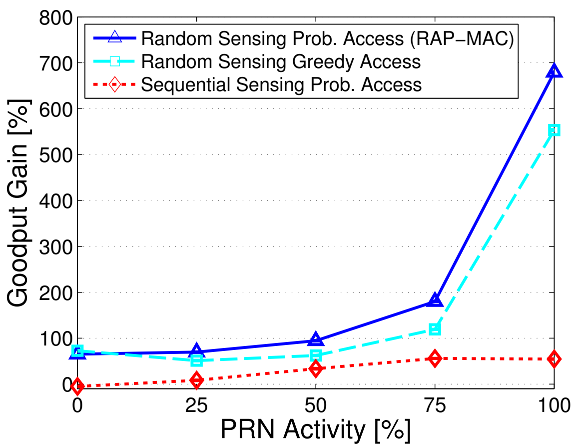

(b) Gain w.r.t. sequential sensing with greedy access.

Fig. 5. RAP-MAC achieves significant goodput gain over traditional opportunistic spectrum access scheme. While both components contribute to the overall gain, the goodput gain due to randomized sensing is higher than the gain due the probabilistic access mechanism.

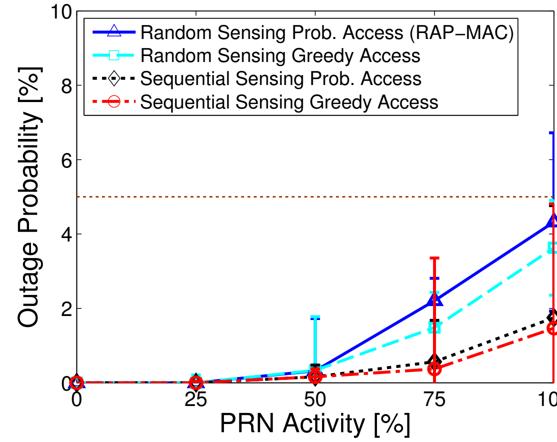

(a) Primary network operating using channel 1.

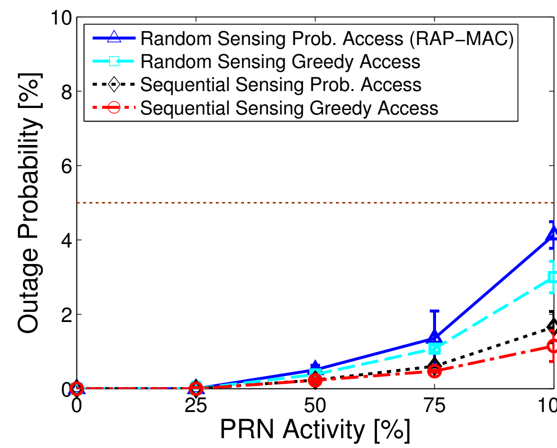

(b) Primary network operating using channel 7.

Fig. 6. The outage probability of the primary networks versus the activity factor for different protocol implementations. While both satisfy the 5% PRN outage constraint, random sensing results in more primary outages compared to sequential sensing.

performance at the expense of having a worse impact on the PRN outages. For a given PRN outage (e.g., 1%), the goodput of RAP (2.55 Mbps) is greater than that of random sensing greedy access (1.76 Mbps) which is greater than the sequential sensing probabilistic access goodput (600 kbps) which is greater than the goodput of sequential sensing greedy access (251 kbps). Meanwhile, for a given goodput, the PRN outages exhibit the opposite trend (i.e., RAP has the hightest outage) as shown in Table II.

C. Performance in Large-Scale Networks

The above experiments demonstrate the superior performance of practical OSA techniques and how traditional OSA approaches can benefit from the gains available to the individual practical components in a simple single-flow topology. Given the WARP cost and complexity of realizing large-scale scenarios, we cannot empirically evaluate the large-scale performance. Here, were present a representative large-scale MATLAB simulation experiment for the sake of completeness. We refer interested readers to [3] for a thorough analytical and simulations-based performance evaluation.

Consider a network setup that is composed of a single CRN collocated with 9 PRNs equally split over the 700 MHz, 2.4 GHZ, and 5 GHz bands. Each PRN and the CRN have 200 users. In each band, we have a lightly-loaded, average-loaded, highly-loaded PRN with respective activities of 10%, 50%, and 90%. The maximum allowed interference by the PRNs is -57 dBm with maximum allowed outage probability, β, of 5% and 10%. We vary the arrival rate of all CRN users from 1 Mbps to 30 Mbps. Figure 7 depicts the RAP goodput gain and fairness gain over the traditional greedy access approaches based on wideband sensing (which resembles sequential sensing benchmark). These results are the average of 30 randomly generated topologies. The fairness gain measures the relative improvement of RAP in Jain's Fairness Index widely used to asses the fair distribution of goodput amongst competing flows. While we could not assess the fairness performance using WARP experiments, Figure 7 shows that RAP achieves a remarkable fairness performance without explicit inter-flow coordination. Furthermore, simulation results exhibited similar increase in the outages experienced by the PRN when RAP is used. However, such simulation-based outages are far below the targeted bound compared to the above WARP empirical

TABLE II
PERCENTAGE OUTAGE OF THE PRNS (OPERATING ON CHANNEL 1 - CHANNEL 7) FOR GIVEN GOODPUT VALUES.

Goodput	Rand. Sens. Prob. Acc.	Rand. Sens. Grdy. Acc.	Seq. Sens. Prob. Acc.	Seq. Sens. Grdy. Acc.
5 Mbps	0.1%-0.1%	0%-0%	0%-0%	0%-0%
4 Mbps	0.36%-0.24%	0.23%-0.18%	0%-0%	0%-0%
3 Mbps	0.51%-0.33%	0.45%-0.3%	0.03%-0.02%	0.01%-0.01%
2 Mbps	1.8%-1.25%	0.85%-0.77%	0.21%-0.16%	0.19%-0.14%
1 Mbps	3.4%-2.9%	2.0%-1.6%	0.39%-0.32%	0.35%-0.28%

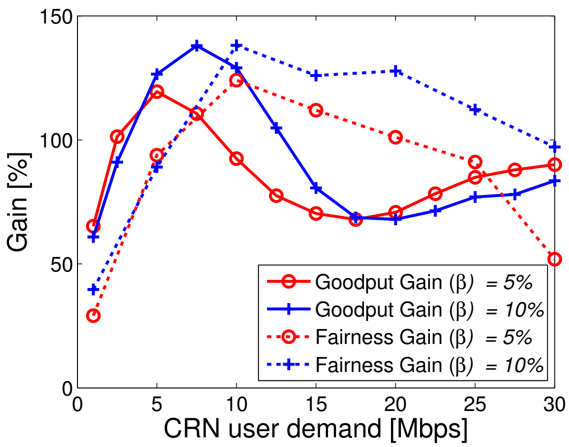

Fig. 7. Simulation results of large-scale networks show that RAP OSA achieves significant goodput and fairness gains with respect to traditional OSA which adopt greedy access based on sensing the entire spectrum for different values of the maximum allowed PRN outage, β.

results. This is due to the underlying idealistic assumptions used in the simulations. Further analysis and simulation results of the PRN outage behavior and its relationship with the RAP probabilistic parameters and the PRN constraints are available in [3] for interested readers.

VI. CONCLUSIONS

In this paper, we have presented an experimental study of the less-well studied topic of distributed opportunistic spectrum access implementation. Our goal has been to demonstrate that while existing hardware technologies do not provide the cognitive transceiver requirements needed to exploit OSA to its full potential, suboptimal OSA approaches developed to target low-complexity transceivers can achieve significant performance improvement compared to theoretically-optimal approaches. More specifically, we have shown that the use of random spectrum selection combined with non-greedy and probabilistic access leads to multiple folds goodput gain at the expense of higher primary outage (within permissible bounds). We have also shown that other theoretically-driven OSA approaches can exploit the gains of either techniques. We plan to further extend our experiments to consider more primary networks and multi-flow cognitive networking to study the fairness performance of different approaches.

REFERENCES

[1] I. F. Akyildiz, W.-Y. Lee, and K. R. Chowdhury, "CRAHNs: cognitive radio ad hoc networks," *Ad Hoc Netw.*, vol. 7, no. 5, pp. 810–836, July 2009.

[2] H. B. Salameh and M. Krunz, "Channel access protocols for multihop opportunistic networks: challenges and recent developments," *IEEE Netw.*, vol. 23, no. 4, pp. 14–19, July-Aug. 2009.

[3] A. Khattab, D. Perkins, and M. A. Bayoumi, "Probabilistic framework for opportunistic spectrum management in cognitive ad-hoc networks," *EURASIP J. Wireless Commun. Netw.*, vol. 2011, no. 188, Nov. 2011.

[4] "Rice University WARP Project." Available: http://warp.rice.edu

[5] "CROSS: Cognitive Radio Open Source System." Available: http://cornet.wireless.vt.edu/trac/wiki/Cross

[6] D. R. DePoy, "Cognitive radio network testbed (cornet): design, deployment, administration and examples," Ph.D. dissertation, Virginia Polytechnic Institute and State University, 2012.

[7] L. Yang, Z. Zhang, W. Hou, B. Y. Zhao, and H. Zheng, "Papyrus: a software platform for distributed dynamic spectrum sharing using SDRs," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 41, pp. 31–37, Jan. 2011.

[8] J. Ansari, X. Zhang, A. Achtzehn, M. Petrova, and P. Mahonen, "A flexible MAC development framework for cognitive radio systems," in *Proc. 2011 IEEE WCNC*.

[9] C. Doerr, M. Neufeld, J. Fifield, T. Weingart, D. Sicker, and D. Grunwald, "MultiMAC—an adaptive MAC framework for dynamic radio networking," in *Proc. 2005 IEEE DySPAN*.

[10] D. G. Messerschmitt, "Rethinking components: from hardware and software to systems," *Proc. IEEE*, vol. 95, pp. 1473–1496, July 2007.

[11] A. Khattab, D. Perkins, and M. A. Bayoumi, "Experimental evaluation of opportunistic spectrum access in distributed cognitive radio networks," in *Proc. 2012 IEEE IWCMC*.

[12] ——, "Opportunistic spectrum access: from theory to practice," *IEEE Veh. Technol. Mag.*, vol. 7, no. 2, pp. 62–68, June 2012.

[13] J. Mitola III, "Cognitive radio: an integrated agent architecture for software defined radio," Ph.D. dissertation, KTH Royal Institute of Technology, 2000.

[14] "Ettus Research LLC." Available: http://www.ettus.com/

[15] "GNU Radio." Available: http://gnuradio.org/trac

[16] K. Tan, J. Zhang, J. Fang, H. Liu, Y. Ye, S. Wang, Y. Zhang, H. Wu, W. Wang, and G. M. Voelker, "Sora: high-performance software radio using general-purpose multi-core processors," *Commun. ACM*, vol. 54, pp. 99–107, Jan. 2011.

[17] K. Tan, H. Shen, J. Zhang, and Y. Zhang, "Enable flexible spectrum access with spectrum virtualization," in *Proc. 2012 IEEE DYSPAN*.

[18] A. Delfino, L. Goratti, R. Giuliani, F. Oliveri, and G. Baldini, "A software radio implementation of centralized MAC protocol for cognitive radio networks," *Wireless Personal Commun.*, vol. 68, no. 3, pp. 1147–1175, 2013.

[19] G. Nychis, T. Hottelier, Z. Yang, S. Seshan, and P. Steenkiste, "Enabling MAC protocol implementations on software-defined radios," in *Proc. 2009 USENIX Symposium on NSDI*.

[20] K. Le, P. Maddala, C. Gutterman, K. Soska, A. Dutta, D. Saha, P. Wolniansky, D. Grunwald, and I. Seskar, "Cognitive radio kit framework: experimental platform for dynamic spectrum research," in *Proc. 2012 ACM WiNTECH*.

[21] A. Young and C. Bostian, "Simple and low-cost platforms for cognitive radio experiments [application notes]," *IEEE Microwave Mag.*, vol. 14, no. 1, pp. 146–157, Jan.-Feb. 2013.

[22] A. Sharma and E. M. Belding, "FreeMAC: framework for multi-channel MAC development on 802.11 hardware," in *Proc. 2008 ACM PRESTO Workshop*.

[23] M.-H. Lu, P. Steenkiste, and T. Chen, "FlexMAC: a wireless protocol development and evaluation platform based on commodity hardware," in *Proc. 2008 ACM WiNTECH*.

[24] A. Sharma, M. Tiwari, and H. Zheng, "MadMAC: building a reconfigurable radio testbed using commodity 802.11 hardware," in *Proc. 2006 IEEE SECON WSDR*.

[25] K. Hong, S. Sengupta, and R. Chandramouli, "Spiderradio: a cognitive radio implementation using IEEE 802.11 components," *IEEE Trans. Mobile Comput.*, vol. 99, Sept. 2012.

[26] Z. Miljanic, I. Seskar, K. Le, and D. Raychaudhuri, "The WINLAB network centric cognitive radio platform—WiNC2R," in *Proc. 2007 CrownComm*.

[27] P. Marshall, "Extending the reach of cognitive radio," *Proc. IEEE*, vol. 97, pp. 612–625, Apr. 2009.

[28] "DARPA's Wieless Network after Next." Available: http://www.darpa.mil/Our_Work/STO/Programs/Wireless_Network_after_Next_(WNAN).aspx

[29] "FCC 10-174: second memorandum opinion and order in ET docket No. 04-186," Sept. 2010.

[30] T. Yucek and H. Arslan, "A survey of spectrum sensing algorithms for cognitive radio applications," *IEEE Commun. Surveys Tuts.*, vol. 11, no. 1, pp. 116–130, First Quarter 2009.

[31] Z. Tian and G. Giannakis, "Compressed sensing for wideband cognitive radios," in *Proc. 2007 IEEE ICASSP*.

[32] H. Liu and B. Krishnamachari, "Randomized strategies for multi-user multi-channel opportunity sensing," in *Proc. 2008 IEEE CCNC Cognitive Radio Netw. Workshop*.

[33] Q. Zhao, L. Tong, A. Swami, and Y. Chen, "Decentralized cognitive MAC for opportunistic spectrum access in ad hoc networks: a POMPD framework," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 3, pp. 589–600, Apr. 2007.

[34] S. Huang, X. Liu, and Z. Ding, "Opportunistic spectrum access in cognitive radio networks," in *Proc. 2008 IEEE INFOCOM*.

[35] F. Wang, M. Krunz, and S. Cui, "Price-based spectrum management in cognitive radio networks," *IEEE J. Sel. Topics Signal Process.*, vol. 2, no. 1, pp. 74–87, Feb. 2008.

[36] H. B. Salameh, M. Krunz, and O. Younis, "MAC protocol for opportunistic cognitive radio networks with soft guarantees," *IEEE Trans. Mobile Comput.*, vol. 8, no. 10, pp. 1339–1352, Oct. 2009.

[37] A. Anandkumar, N. Michael, and A. Tang, "Opportunistic spectrum access with multiple users: learning under competition," in *Proc. 2010 IEEE INFOCOM*.

[38] P. Chaporkar, A. Proutiere, and H. Asnani, "Learning to optimally exploit multi-channel diversity in wireless systems," in *Proc. 2010 IEEE INFOCOM*.

[39] J. Jia, Q. Zhang, and X. Shen, "HC-MAC: a hardware-constrained cognitive MAC for efficient spectrum management," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 1, pp. 106–117, Jan. 2008.

[40] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "Predictable 802.11 packet delivery from wireless channel measurements," in *Proc. 2010 ACM SIGCOMM*.

[41] "RAP-MAC demo." Available: http://www.youtube.com/watch?v=8X-cbQFrLxM

Ahmed Khattab received his B.Sc. (honors) and M.Sc. in Electronics and Communications Engineering from Cairo University, Egypt, in 2002 and 2004, respectively. He received the Master of Electrical Engineering degree from Rice University, and his Ph.D. in Computer Engineering from the University of Louisiana at Lafayette, USA, in 2009 and 2011, respectively. He is currently an Assistant Professor at the EECE Department in Cairo University. His research interests are in wireless communication networks and embedded systems. He won the best student paper award of the UL IEEE Computer Society in 2010 and 2011 and was a finalist in the IEEE ICCCN 2008 best paper contest.

Dmitri Perkins received his B.S. degree in Computer Science in 1995 from the Tuskegee University and his M.S. and Ph.D. degrees in Computer Science from Michigan State University in 1997 and 2002, respectively. He is currently an Associate Professor at The Center for Advanced Computer Studies (CACS), University of Louisiana at Lafayette. His primary research interests include performance analysis, scalable protocol design, resource allocation, and self-management issues in multi-hop wireless broadband networks, sensor ad hoc networks and multi-hop cognitive radio networks. Dr. Perkins received the NSF CAREER Award in 2005 and was awarded the Hardy Edmiston Endowed Professorship in 2008.

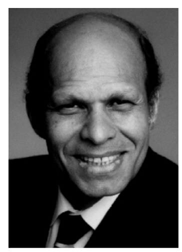

Magdy A. Bayoumi received the BSc and MSc degrees in electrical engineering from Cairo University, Egypt, the MSc degree in computer engineering from Washington University, St. Louis, and the PhD degree in electrical engineering from the University of Windsor, Ontario. He is the Z.L. Loflin Eminent Scholar Endowed Chair at the Center for Advanced Computer Studies (CACS), University of Louisiana, Lafayette. He is a fellow of the IEEE. He was the Vice President for Conferences of the IEEE Circuits and Systems (CAS) Society. He is the recipient of the 2009 IEEE Circuits and Systems Meritorious Service Award and the IEEE Circuits and Systems Society 2003 Education Award. Dr. Bayoumi research interests include VLSI design and architectures, digital signal processing, and wireless ad-hoc and sensor networks.