C-MAC: A Cognitive MAC Protocol for Multi-Channel Wireless Networks

Carlos Cordeiro and Kiran Challapali
Philips Research North America
Briarcliff Manor, NY 10510
{Carlos.Cordeiro, Kiran.Challapali}@philips.com

Abstract – **A number of algorithmic and protocol assumptions taken for granted in the design of existing wireless communication technologies need to be revisited in extending their scope to the new cognitive radio (CR) paradigm. The fact that channel availability can rapidly change over time and the need for coordinated quiet periods in order to quickly and robustly detect the presence of incumbents, are just some of the examples of the unique challenges in protocol and algorithm design for CR networks and, in particular, in the medium access control (MAC) layer. With this in mind, in this paper we introduce a novel cognitive MAC (C-MAC) protocol for distributed multi-channel wireless networks. C-MAC operates over multiple channels, and hence is able to effectively deal with, among other things, the dynamics of resource availability due to primary users and mitigate the effects of distributed quiet periods utilized for primary user signal detection. In C-MAC, each channel is logically divided into recurring superframes which, in turn, include a slotted beaconing period (BP) where nodes exchange information and negotiate channel usage. Each node transmits a beacon in a designated beacon slot during the BP, which helps in dealing with hidden nodes, medium reservations, and mobility. For coordination amongst nodes in different channels, a rendezvous channel (RC) is employed that is decided dynamically and in a totally distributed fashion. Among other things, the RC is used to support network-wide multicast and broadcast which are often neglected in existing multi-channel MAC protocols. We present promising performance results of C-MAC. We also describe our efforts to implement features of C-MAC in a real CR prototype with Atheros chipset, which currently includes the spectrum sensing module and preliminary features of C-MAC.**

I. INTRODUCTION

Cognitive radios (CRs) [1][2][3] have emerged as the solution to the problem of spectrum scarcity for wireless applications. It is the key technology that will enable flexible, efficient and reliable spectrum use by adapting the radio's operating characteristics to the real-time conditions of the environment. CRs have the potential to utilize the large amount of unused spectrum in an intelligent way while not interfering with other incumbent devices in frequency bands already licensed for specific uses.

In terms of functionality, the key goals of a CR are to enable dynamic spectrum access (DSA), dynamic spectrum sharing (DSS), and dynamic spectrum multi-channel (DSM) [4][5]. A critical ingredient to meet these goals is the medium access control (MAC) layer, which is responsible for many of the radio control functions such as those for timely and reliable incumbent detection (i.e., DSA), resource sharing amongst networks (i.e., DSS) and frequency agility (i.e., DSM).

Therefore, in this paper we introduce a novel cognitive MAC (C-MAC) protocol for decentralized multi-channel CR networks. It is well known that by simultaneously operating over multiple channels (hence requiring a multi-channel MAC) the capacity of the wireless network is significantly increased, even if each node only occupies one channel at a time [6]. Decentralized cognitive radio networks are inherently multi-channel in nature, since these networks operate over a set of channels whose availability changes over time depending upon the incumbent duty cycle. To seamlessly vacate a channel and operate on another channel without disruption of services to higher layer, these secondary CRs must monitor other channels as they communicate on a given channel.

The C-MAC protocol employs the concept of dynamic rendezvous channel (RC), which is used to coordinate nodes in different channels, for multi-channel resource reservation, quiet period (QP) coordination for incumbent detection, and so on. Out of all channels available, the RC is assigned the most reliable of them (in terms of availability). Given that incumbents may appear at any time, the concept of Backup Channel (BC) is introduced and is employed to make the RC extremely robust to incumbents. The BC is determined by out-of-band measurements carried out by nodes whenever they are not engaged in communications or, as a last resort, during QPs.

A number of key challenges are addressed by C-MAC for proper operation in the presence of incumbents. To that end, coexistence mechanisms are one of the design cornerstones of C-MAC. Nodes perform in-band and out-of-band measurements via quiet periods to detect incumbents at low signal-to-noise ratio (SNR) and to determine a suitable BC for a given channel. A dynamic inter-channel coordination scheme allows multiple vacant channels to be simultaneously exploited, while offering resilience against incumbents. Multicast and broadcast capabilities are supported through the RC, which are often neglected features in existing multi-channel MAC protocols. The multi-channel hidden terminal problem [16] is overcome by a distributed beaconing approach which synchronizes nodes in time, space, and frequency. For that, each channel is logically divided into recurring superframes that begin with a slotted beacon period (BP) followed by a data transfer period (DTP). During the BP, each node transmits a beacon in its designated time slot. Beacons contain information about scheduled QPs, spectrum

1-4244-0663-3/07/$20.00 ©2007 IEEE

measurements, and multi-channel reservation for data communication. Once the BP is over, nodes may switch to other channels for communication.

To validate the design of C-MAC, we have embarked on analytical, simulation and prototyping efforts. Here we present some preliminary performance results of C-MAC including a description of our prototyping efforts and outcomes.

The rest of this paper is organized as follows. Section II presents the related work in the area of MAC protocols for CRs. Next, Section III discusses some open issues in the area and motivates the need for C-MAC. This is followed by Section IV which presents the proposed C-MAC protocol and how it addresses the challenges posed by the CR paradigm. Performance evaluation results are then described in Section V. Finally, this paper is concluded in Section VI.

II. RELATED WORK

As far as MAC layer for CRs goes, research is still in its infancy. The IEEE 802.22 working group is in the process of standardizing a MAC layer based on CR for reuse of spectrum that is allocated to TV broadcast service [7]. The architecture of the 802.22 MAC layer is centralized and relies on a base station, while C-MAC is designed for a fully decentralized operation.

In [8] a very high level overview of a DSA system is presented, with little or no details given as to the algorithms and protocols used. The Dynamic Open Spectrum Sharing (DOSS) MAC protocol is introduced in [9], and is a multi-channel MAC that incorporates the busy-tone concept to overcome the hidden and exposed node problem in wireless networks. While DOSS allows nodes to dynamically negotiate the channel to be used for data communication based on spectrum availability, it does not address all the critical aspects related to CR operation, such as sensing algorithms, dynamic device discovery without a fixed control channel, network recovery, and so on. DOSS also requires multiple radio transceivers. In [10] it provided a theoretical formulation of a decentralized MAC, with no insights into protocol design, implementation and performance.

Although not designed specifically for CR networks, there have been a number of recent proposals addressing the problem of coordinated use of multiple channels at the MAC layer (i.e., multi-channel MAC). All these protocols, however, have similar limitations and do not deal with the new challenges posed by CRs. Nevertheless, for completeness purposes it is important to provide a description of their operation. These multi-channel MAC protocols can be classified based on how many radio transceivers they require for operation, namely single transceiver protocols or multiple transceiver protocols.

A. Single Transceiver Protocols

This category of MAC protocols assume that every node is equipped with one half-duplex transceiver capable of switching channels dynamically, and it can only transmit or receive on exactly one channel at any given time. Protocols in this category often aim at incurring a complexity comparable to existing solutions (e.g., IEEE 802.11), while achieving better throughput and delay performance. Some protocol design challenges are how to overcome the hidden and exposed terminal problem with low control overhead, minimize channel switching, load balancing, achieve network connectivity comparable to single channel MAC protocols (e.g., IEEE 802.11), and so on.

The Hop Reservation Multiple Access (HRMA) protocol [11] is a multi-channel MAC scheme for slow frequency hopping spread spectrum (FHSS) wireless ad hoc networks where all nodes hop according to a pre-defined hopping pattern. Whenever a node has a data packet to send, it exchanges RTS/CTS packets with the intended receiver and both remain in the same hop for the entire data transmission. Other nodes not involved in communication do not stop and proceed by following the hopping sequence. Since different pairs of nodes can communicate simultaneously while in different hops, HRMA is considered as a multi-channel MAC. While in HRMA it is the sender node who initiates communication, in Receiver Initiated Channel-Hopping with Dual Polling (RICHDP) protocol [12] this responsibility is transferred to the receiver. Other than this, HRMA and RICHDP behave similarly. Since these protocols have been designed for FHHS, they cannot be applied to the popular Direct Sequence Spread Spectrum (DSSS) systems under consideration.

In [13] it is considered that the number of nodes in the network equals the number of available channels. Out of the total N channels, one is reserved as a default control channel while the others are employed for data transmissions. Before any data communication, the sender node has to negotiate with the receiver a data channel through a RTS/CTS handshake transmitted in the control channel.

In [14], every node is associated with a single channel which is derived on the basis of a node's MAC address. This particular channel is referred to as home channel and is used by the node to wait for incoming packets. A node S wishing to communicate to a node D would have to switch to node D's home channel before transmission, and immediately return to its home channel after completion.

The Channel Hopping Multiple Access (CHMA) [23] and the Slotted Seeded Channel Hopping (SSCH) algorithm [24] use a similar channel hopping approach (with some variation on the hopping pattern generation). If a node wants to communicate with another node, it follows the other node's schedule. If two nodes are able to successfully exchange control information, they stay on that channel to complete the data transfer.

Switching amongst channels may take considerable time and hence may increase delay and degrade throughput. With this in mind, the On-Demand Channel switching (ODC) [15] mechanism aims at minimizing such negative impact by having nodes stay in its channel as long as traffic conditions on this channel are acceptable. Nodes continuously measure channel conditions and use this measurements for switching decision. As all channels are equal in ODC, finding intended receivers is more difficult. In addition, ODC performance is not uniform and is very dependent on the traffic pattern.

The Multi-channel MAC (MMAC) [16] protocol has the primary goal of overcoming the multi-channel hidden terminal problem present in many multi-channel MAC protocols based on a single transceiver. It reuses the Power Saving Mode (PSM) concept of IEEE 802.11 and its corresponding Ad-Hoc Traffic Indication Messages (ATIM) control messages. On the basis of this, it defines a default control channel where all nodes must periodically switch to and synchronize for a pre-determined window of time. This is called the ATIM window and where nodes with packets to send employ a three-way handshake (ATIM/ATIM-ACK/ATIM-RES) as to negotiate a data channel. Communicating nodes may then switch to the selected channel and contend for medium access by using traditional RTS/CTS/Data/ACK mechanism.

The Multi-channel Access Protocol (MAP) [22] is based on a similar concept as MMAC, and divides the time into control periods, during which all the nodes tune to the control channel for control message exchanging, and data periods, during which data transfer takes place.

B. Multiple Transceiver Protocols

When multiple transceivers are in place, the task of designing a multi-channel MAC protocols is significantly simplified. Issues such as hidden and exposed terminal problems, connectivity, and channel switching can be overcome almost completely. Here, it is assumed that nodes have multiple half-duplex transceivers capable of tuning to and accessing different channels simultaneously, which is the key to overcoming the aforementioned challenges. Research here has mostly focused on channel selection strategies.

In [17] it is introduced the Dynamic Private Channel (DPC) protocol where nodes are assumed to be equipped with as many transceivers as the number of channels. Similar to other protocols, one particular channel is reserved as the default control channel for negotiation purposes. Given that a transceiver is always associated with the control channel, the multi-channel hidden terminal problem is eliminated. Special RTS and reply-to-RTS packets are employed in this control channel in order to select another traffic channel for data communication. Once the traffic channel is negotiated, nodes exchange CTS/Data/ACK packets through the transceiver associated with the selected channel.

The multi-channel MAC protocol proposed in [21] also assumes that each node has as many transceivers as there are channels, but here nodes are capable of listening to all these channels simultaneously. Whenever a node has a packet to send, it selects an idle channel for transmission. In case of multiple idle channels, the one employed in the last successful data transmission is preferred. This technique is referred to as “soft channel reservation”. An enhanced channel selection strategy for this protocol has been presented in [18] and consists in selecting the best channel based on the power level sensed at the transmitter.

On the other hand, the Receiver-Based Channel Selection (RBCS) mechanism in [19] chooses the best channel on the basis of the signal-to-interference and noise ratio (SINR) at the receiver. To this end, RTS/CTS packets are employed in a default control channel as to select the data channel with highest SINR.

The Dynamic Channel Assignment (DCA) protocol [20] operates similar to RBCS. It employs a default control channel while other channels may be used for data transmission. RTS/CTS packets are exchanged in the control channel and serve to negotiate a data channel for Data/ACK transmission. A distinctive feature of DCA is that is requires exactly two transceivers, one of which is permanently tune to the default control channel and the other which is free to tune to any of the data channels. As noted in [16], a drawback in DCA is that it dedicates one channel for exchanging control information only. When the number of channels is small (e.g., only 3 channels in IEEE 802.11b), this constitutes a considerable wastage of resources.

Based on MMAC [16], the power saving multi-channel MAC protocol (PSM-MMAC) [28] targets to reduce power consumption under multi-channel operation, which is highly desired due to the fact that some nodes are powered by battery. Moreover, PSM-MMAC facilitates heterogeneous nodes to coordinate with others in such a way that powerful nodes (with multiple transceivers) can take advantage of more channels if desired. However, PSM-MMAC focuses only on the one-hop case. It is not straightforward to apply it directly to the multi-hop case.

Finally, the Common Spectrum Coordination Channel (CSCC) protocol [25] is an extension of the DCA protocol that allows different types of wireless devices to share the radio spectrum. This is done via negotiation through the CSCC.

III. ISSUES AND MOTIVATION

The benefits brought by multi-channel operation in CR networks, such as higher aggregate throughput and better robustness, should not come at the expense of other factors such as increased cost and complexity. With this in mind, some of the key requirements for practical multi-channel wireless MAC protocol are negligible performance disruption due to shared operation with incumbents, complexity and network connectivity comparable to those provided by single channel networks (e.g., those based on IEEE 802.11), multi-channel hidden terminal problem, load balancing, efficient radio resource usage, group communication support (broadcast and multicast), and mobility support.

Even though not designed for CR networks, MMAC is one of the most prominent multi-channel protocols that attempt to address some of these challenges with a single transceiver. It proposes a solution to the hidden terminal problem when all devices are within radio range of each other and provides some level of load balancing. Despite of that, MMAC fails to tackle these aspects satisfactorily while leaving open many important issues.

In particular, the multi-channel hidden terminal problem is a major challenge in multi-channel MAC protocols based on a single radio transceiver. It occurs when a node A is tuned to one particular channel and hence cannot listen to any other channel. Here, node A is *hidden* from any ongoing conversation in another channel. In case node A dynamically switches to a busy channel and transmits, a collision may take place. Clearly, if at least two transceivers are available this problem can be easily overcome by having one of the transceivers permanently tuned to a pre-defined RC, while the other would be free to switch amongst data channels.

MMAC overcomes the multi-channel hidden terminal problem and provides connectivity and complexity comparable to single channel networks, but only when all nodes are within radio range of each other. In other words, multi-hop scenarios are not handled by MMAC. This is due to the fact that in MMAC only those nodes which are engaged in communication actually exchange ATIM control messages. In a multi-hop scenario where not all transmissions are heard, MMAC fails as exemplified in Figure 1(b) for the multi-hop topology of Figure 1(a). In this figure, nodes A and B continuously perform the three-way handshake through ATIM packets while nodes C and D remain silent as they do not have packets to be transmitted. Given that C and D are not participating in the ATIM exchange, after some time node D's clock (and perhaps node C) will drift away and hence will lose synchronization with the rest of the network. Once this happens, MMAC will start suffering from the hidden terminal problem and its negative impacts in case, for example, node D decides to communicate with node C.

In case we add mobility to this scenario, MMAC performance can be further impacted. For example, consider in Figure 1(a) that nodes A and B are communicating in channel 1 when node D moves within node B's radio range. It may occur that nodes C and D exchange ATIM packets and select the same channel 1 where nodes A and B are currently communicating. This is possible as the beacon intervals of nodes A and B, and nodes C and D are shifted in time. As a result, the transmission between these two pair of nodes may collide.

Another major drawback in MMAC is depicted in Figure 2, wherein a total of three channels are available and with channel 0 as the default control channel where ATIM packets are exchanged. As we can see from this figure, there is a major channel bandwidth wastage in channels 1 and 2 for the time during which ATIM packets are exchanged in the default channel 0 (the ATIM window usually occupies as much as 20% of the superframe). If the number of channels a large, this wastage will clearly increase thus making MMAC extremely bandwidth inefficient.

MMAC also requires a high overhead for achieving its goals. For any data packet to be sent, a pair of nodes has to perform a three-way handshake in the control channel and a RTS/CTS handshake in the data channel. Moreover, these handshakes are all preceded by a backoff phase. While this may be reasonable in low load environments, in medium to high load scenarios this overhead may become unacceptable.

Support for group communication (GC) mechanisms such as broadcasting and multicast is another shortcoming of mostly all multi-channel MAC protocols. If we take MMAC as an example, we can clearly see that there is no way a device is able to transmit a broadcast/multicast data packet given that it relies on ATIM control packets for channel negotiation.

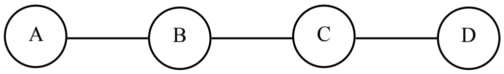

(a) – Network Topology

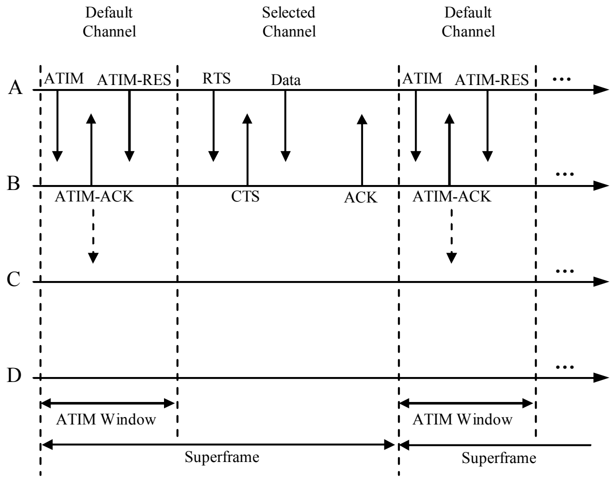

(b) – Nodes A and B negotiate a channel for data exchange

Figure 1 – MMAC protocol operation when nodes A and B communicate. If nodes C and D do not initiate communication within a certain amount of time, they may become desynchronized with the rest of the network. This may lead to the reappearance of the multi-channel hidden terminal problem.

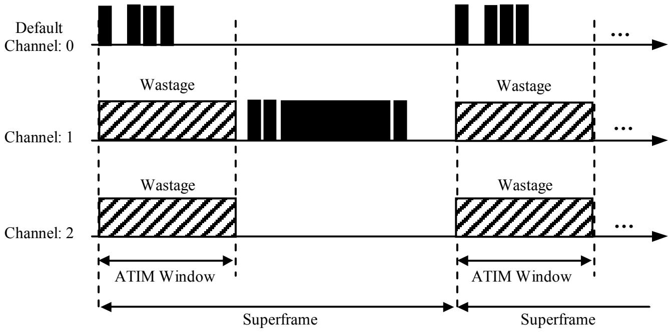

Figure 2 – In MMAC, whenever ATIM packets are exchanged in the default common channel valuable bandwidth is wasted in all other channels. As the number of channels increase, this wastage will also increase.

Finally, load balancing in MMAC is accomplished by counting the number of ATIM-ACK or ATIM-RES packets a node overhears (i.e., counting the number of communicating node pairs). A channel with the lowest value for this counter is then selected for further data transmission. A clear limitation of this mechanism is when nodes have data packets of various lengths. In this case, counting the number of communicating node pairs will mislead the protocol and provide very poor accuracy for the channel selection procedure.

IV. THE PROPOSED C-MAC PROTOCOL

The C-MAC protocol has been designed to address the requirement of decentralized and distributed operation on a non-interfering basis. It not only incorporates novel mechanisms dealing with multi-channel operation, but also does this in conjunction with robust incumbent protection mechanisms.

A. Assumptions

To keep the complexity comparable to existing wireless standards while fully exploiting the flexibility of multiple channels, C-MAC assumes that terminals are equipped with a single half-duplex radio. In other words, at any given time a host is capable of either transmitting or receiving, but not both. In addition, a terminal can only receive or send in one channel at a time, in such a way that when the terminal is receiving in one channel it cannot perform carrier sense in another.

B. Multi-Channel and Superframe Structures

While there is plenty of research in channel structure for single channel MAC, to the best of our knowledge no existing work has addressed this issue from a multi-channel perspective.

To fully exploit the multiple channels available in a scalable manner, C-MAC incorporates the channel structure illustrated in Figure 3. Each channel has its own superframe structure, and out of all the channels in use one is uniquely identified as the RC (channel A in Figure 3 – see also Section IV.C). This is in contrast to the channel structure employed in MMAC where a superframe is only used in the common channel, hence requiring all network devices to switch back to the common channel upon the start of every new superframe (which limits scalability). Another added benefit of the channel structured adopted in C-MAC is that it also overcomes the bandwidth wastage of MMAC, as beacons can be transmitted on any channel and not only on the common channel. Hence, this leads to a better load balancing and allows more time for actual data transfer.

Each superframe is comprised of two consecutive parts: the Beacon Period (BP) and the Data Transfer Period (DTP). A distinctive feature of this multi-channel structure is that the BPs across channels are non-overlapping (see Figure 3). This is done through the inter-channel synchronization mechanism described in Section IV.E, and allows device to quickly gather information about other channels in an optimized fashion by simply switching channels in ascending order of BP start time (BPST) and listening for beacon frames during the BP. If, on the other hand, all a device needs is information about which node is located on which channel, this can be efficiently obtained from beacons received on the RC channel itself, as also described in Section IV.E.

Within a superframe, C-MAC employs a slotted access mechanism as presented in Figure 4. The start time of a superframe is determined by the beginning of the beacon period (i.e., BPST). The first two slots of the BP are termed as the signaling slots, and are used for new devices joining this channel. From the third slot onwards, each device on that channel sends its own beacon during the BP at its designated beacon slot, as shown in Figure 4. Since terminals participating in the superframe typically use the same BPST, collisions are not a problem as slots are indexed relative to the same BPST.

C. The Rendezvous Channel

The RC can be seen as the backbone of C-MAC. It is used as the means to manage the entire network. Through the RC, the following features can be accomplished:

- Network-wide GC: The RC is used as the vehicle to send GC packets across channels.
- Inter-channel synchronization: Nodes that do not send beacons within the RC periodically visit the RC to get resynchronized. This is used later on to adjust the BPST of all channels as to make them non-overlapping.
- Neighborhood discovery: since nodes have to periodically visit the RC for resynchronization, they can quickly figure out information about network connectivity.
- Load balancing: the RC acts a conduit for sharing load balancing information of different channels. The channel selection algorithm takes this into account for balancing load.

As we can see, the selection of the RC is a critical component of C-MAC and works as follows. Upon power up, each device scans all the available channels performing measurements and also looking for any beacons frames transmitted by other devices. In those channels where incumbent signals are not sensed, the device dwells on the channel for at least one superframe length as to guarantee that it sees a beacon frame. In case a device receives one or more beacon frames on one channel, it reads the RC field available from the beacon frame header. If this bit is set (meaning this is a RC), the device may decide to join this BP by sending its own beacon during the signaling slot and then move to a permanent and designated beacon slot. Otherwise, the device may decide to continue the scan procedure looking for a RC. If the device scans all channels without detecting any beacon frame with the RC field set, it will itself select a channel as a RC and will start transmitting beacons with the RC field set to one. On the other hand, a device cannot initiate a new RC if it has detected the existence of a RC.

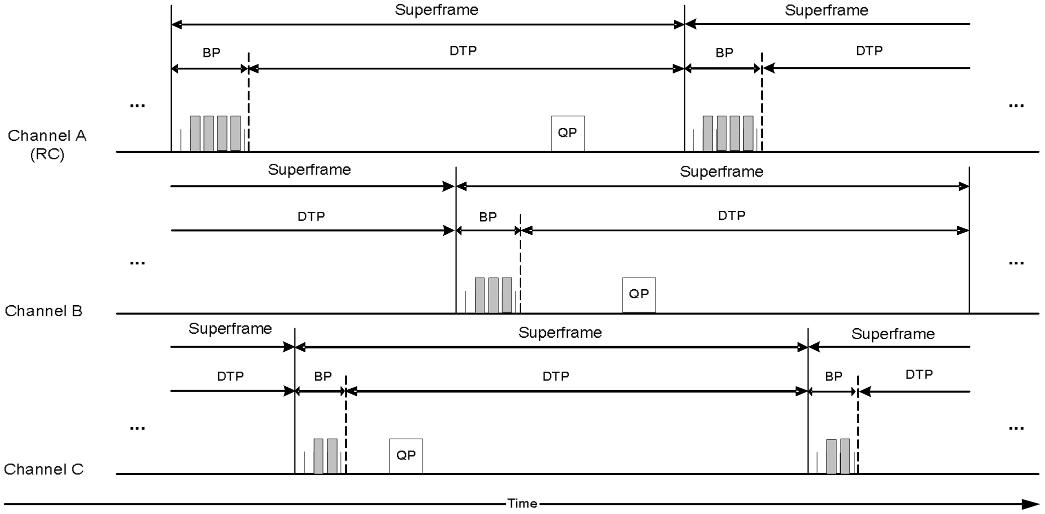

Figure 3 – Multi-channel superframe structure in C-MAC. Here, each channel is structured in the form of superframes whose BPs are non-overlapping across channels.

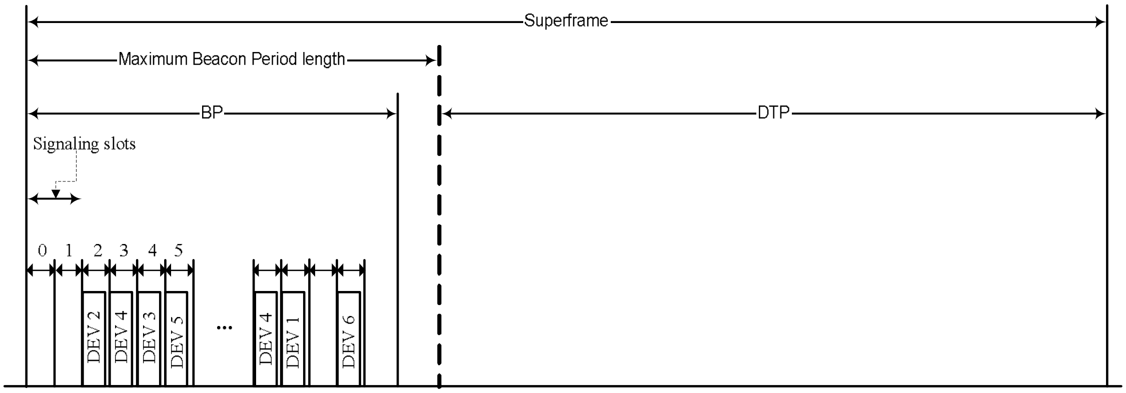

Figure 4 – Slotted access structure of a superframe

It is possible that there exist more than one RC before they eventually converge. Therefore, even after devices are already associated with one RC, they periodically scan other available channels. This is called out-of-band measurements, which are used to, among other things, detect the presence of incumbents, identify other overlapping RCs, determine a suitable BC, and collect channel quality information (see Section IV.G for more on out-of-band measurements). If during this procedure the device detects the presence of other RCs, it will initiate RC switch as follows. Once this device returns to its primary RC and transmits its regular beacon frame during the BP, it includes a RC-switch information element in the regular beacon. This element indicates the new RC as well as schedules the RC switch to some random time in the future. All devices in the primary RC record this information and once this timer expires, they all switch as a group to the new RC. Here, we note that since the RC switch is sufficiently random, it is very unlikely that different set of devices in different RCs decide to switch simultaneously. Therefore, after a finite amount of time all devices in the same neighborhood will converge to the same RC. Once the devices switch to the new RC, they start sending beacons which avoids any instability situation where groups of devices end up switching channels back and forth.

D. Distributed Beaconing

In C-MAC, every terminal is required to transmit a beacon during the BP of a superframe. In the beacon that it transmits, the device rebroadcast information that it received from its neighbors in the previous superframe. Thus, devices have the information about their neighbor's neighbors, such as occupied beacon slot and communication schedules. With this mechanism, it is possible to support mobility and overcome the multi-channel hidden terminal problem.

To understand how this is accomplished, we introduce the concepts of beacon group and extended beacon group, which are depicted in Figure 5. Beacon group and extended beacon group are defined as to allow contention-free frame exchanges while exploring spatial reuse. A beacon group is defined as the set of devices from which a device receives beacons and that have the same BPST as the receiving device. The extended beacon group is the union of a device's beacon group and the beacon groups of all devices in the device's beacon group. Figure 5 shows the beacon group of devices A, B, C, and D, which is formed by the direct neighbors of a device. The extended beacon group of B is formed by the union of the beacon groups of nodes A, B, and C, while device D does not belong to the extended beacon group of B.

In protocol terminology, a beacon group of a device A is formed by all the devices from which A receives a beacon from. The extended beacon group of A is the union of the beacon group of A and the neighbors' neighbors, which is obtained from the received beacons.

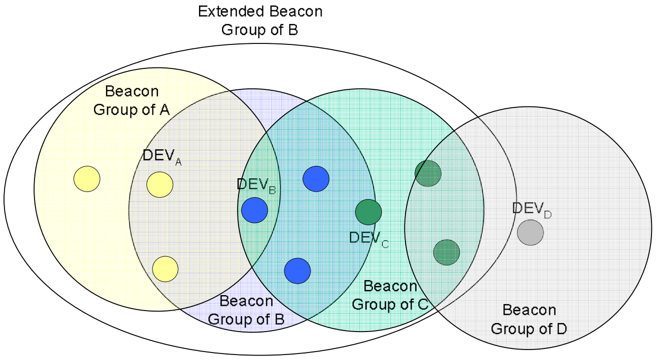

Figure 5 – The concept of beacon group and extended beacon group

With this distributed beaconing procedure where every device sends a beacon, mobility can be better supported. This is due to the fact that even as devices move, they keep sending beacons and so tracking is easier. Also, the hidden terminal problem is addressed, since nodes know about their neighbors' neighbors and about their data transmission schedules. Moreover, since each node sends a beacon in every superframe, periodic resynchronization is accomplished and the clock drifting problem is overcome.

Nevertheless, the distributed beaconing only solves part of the multi-channel hidden terminal problem. As described in Section III, another instance of this problem occurs given the fact that nodes have no knowledge of what is going on in other channels. To address this, we use the RC and enforce protocol rules on nodes. This is described in further details in the next subsection.

E. Inter-Channel Coordination

Inter-channel coordination is an important feature of C-MAC, as it allows nodes to distribute themselves across channels in a coordinated fashion. As described earlier, the core foundation of C-MAC is the RC. Nodes have to first establish a RC before engaging in any sort of communication. Once the RC is setup, nodes can then exploit other channels not occupied by incumbents.

To do it so, nodes have to first announce their impending channel change via a beacon transmission over the RC (see Section IV.F for information on how channels are selected). This is done by appending a channel switch information element to the node's beacon transmission. Other nodes on the RC receiving such beacon record that the transmitting node is going to leave the RC to another channel and rebroadcast it in their beacon. Once the channel changing node hops to its new channel, it sets up the channel by transmitting its own beacon and hence starting a new sequence of recurring superframes. Other nodes in the RC that want to switch to the same channel follow the same procedure, with the difference that they join an existing superframe instead of starting a new one. Figure 3 depicts one example of this procedure. Note that the requirement that a device must always send a beacon before engaging on any data communication on a channel effectively overcomes the multi-channel hidden terminal problem, both in the single-hop as well as in the multi-hop case.

Once devices move out of the RC and set up other channels for the purpose of data communication and load balancing, clock drifting will cause them to lose synchronization with the RC. Also, due to mobility, the network connectivity information which is all managed within the RC may quickly become obsolete.

To address this, devices which operate on other channels other than the RC are required to periodically switch back to the RC as to resynchronize. Resynchronization here means not only timing but also connectivity, and is both achieved by having the device transmit a beacon on the RC upon switching back to the RC for resynchronization. Timing resynchronization allows BPs and QPs of different channels to be kept non-overlapping as shown in Figure 3. This is done by dynamically adjusting the superframe's BPST as devices resynchronize with the RC. Connectivity resynchronization addresses the mobility aspect by keeping the information in the RC up-to-date. Given that a device transmits a beacon every time it tunes back to the RC, the connectivity information can be immediately updated. Once resynchronized, devices may leave the RC for a number of superframes (which, among other things, depends on the relative clock drift amongst devices) until they have to again regain synchronization.

After devices move out of the RC to another channel, say, channel C, all data communication of those devices must happen on channel C. If the devices on channel C want to switch to other channels other than the RC, they have to first switch back to the RC and follow the same procedure as described above. Among other things, this is required as to maintain the information on the RC accurate throughout the network lifetime.

F. Channel Selection and Load Balancing

Contrary to a number of existing multi-channel MAC protocols that include complex channel selection and negotiation schemes, this is not needed in C-MAC. Once on a channel, all data communication of a node happens on that channel. If a node wants to communicate in some other channel which is not the RC, it has to first visit the RC first as described in Section IV.E.

Because of this design choice, however, load balancing becomes a more critical issue in C-MAC. It is desirable to even the traffic load across channels, so that their utilization and the overall network performance are maximized.

The first aspect to load balancing is determining the load of each individual channel. In C-MAC, nodes can easily figure out the load on a particular channel by analyzing the beacon frames transmitted during a BP. This is possible due to the fact that beacons carry the transmitting node traffic reservations for the current superframe. Hence, once a node receives all beacon frames transmitted during a BP, it can deduce the overall channel load in a straightforward manner. At the time the node goes back to the RC to regain synchronization, it also advertises this channel load. Therefore, nodes can decide to which channels to switch to based on the load statistics of that channel. They can also decide to change their current channel.

Obviously, another contributing factor to how a node selects a channel is on which channel the destination of a node is. However, if we consider a uniform traffic pattern across the network, this load balancing mechanism via the RC can effectively balance the load and optimize the aggregate performance.

G. Group Communication

Ideally, a multi-channel MAC protocol should support the same level of connectivity of single-channel protocols, but with higher performance. Towards this goal of providing seamless connectivity, the support to GC (such as multicasting and broadcasting) is critical. Without it, higher layers cannot operate properly as many mechanisms such as address resolution cannot be carried out.

While the task of GC is easily accomplished in protocols such as IEEE 802.11 given its single channel operation, it becomes a hard issue in a multi-channel MAC protocols. Even though devices may be neighbors in physical terms, they may not be capable of hearing each other if they belong to different channels. As a result, GC in a multi-channel environment needs special care even though it has been completely neglected in existing protocols.

We address this problem in C-MAC by employing the RC. Whenever a device S has a GC packet to send, it tunes to the RC and transmits its beacon with the destination address set to the corresponding GC address (i.e., multicast or broadcast address) and with a scheduled transmission time for the GC packet. Node S does this for several consecutive superframes. Devices tuned to the RC and that receive the beacon from S realize that a GC transmission is forthcoming and, if desired, scheduled for the reception of the GC packet at the expected transmission time. Since devices typically have to periodically tune back to the RC for resynchronization, they will eventually receive the beacon from node S notifying of the scheduled GC packet transmission.

With this simple mechanism, C-MAC is able to effectively support GC across multiple channels by means of the RC. It is important to note that, in the case of multicast [26], only those devices participating in the multicast group would schedule the reception of the GC packet in the RC. This is done by having devices inspect the destination address of the GC packet before scheduling the packet reception.

G. Coexistence

Effective coexistence is one of the key responsibilities features of C-MAC, where it takes two forms: i) coexistence with incumbents; and ii) self-coexistence. As the name suggests, coexistence with incumbents deals with DSA mechanisms for a reliable, efficient, and timely detection of primary services (alternatively, detection of white spaces), followed by a network recovery procedure once these incumbents are detected. In contrast, self-coexistence addresses DSS amongst collocated networks operating under C-MAC.

The measurement capability of C-MAC is similar to what is specified in the IEEE 802.22 draft standard [7]. It includes a comprehensive measurement and spectrum management component that provides the necessary flexibility and efficiency. Through beacons, nodes negotiate to perform periodic measurement activities, which may be either in-band or out-of-band. In-band measurement relates to the channel(s) used by the node for communication, while out-of-band correspond to all other channels.

G.1 Coexistence with Incumbents

Coexistence with incumbents is a multi-stage process and involves detection, notification and recovery.

G.1.1 Incumbent Detection

For in-band measurements QPs are used so that incumbent detection through spectrum sensing can be carried out, which is not the case for out-of-band measurements. Therefore, in-band sensing requires tight control of C-MAC, while out-of-band sensing is less critical. C-MAC enforces that QPs across channels be non-overlapping as shown in Figure 3. By doing that, devices can schedule their out-of-band sensing of other channels during that channel's QP. For example in Figure 3, devices operating on channels B and C can perform their out-of-band sensing of channel A during channel A's QP. For that, QP schedules of different channels are also communicated through a device's beacon once it goes back to the RC for resynchronization. As for specific spectrum sensing algorithms used in C-MAC, please see [4].

G.1.2 Incumbent Notification

Once an incumbent is detected, this event must be notified in a timely fashion to other network devices. In C-MAC, this is done through the beacon frame sent by devices in every superframe. Since the beacon frames are sent at the most robust modulation and coding, it can be received despite the presence of the incumbent signal. It is also possible, however, that the incumbent signal power level is so high (e.g., a TV station) that no beacons can be received. In this case, however, there is a high probability that if one device detects the presence of a high power signal, so will its neighbors. Thus, this procedure is facilitated. As a last resort, timeouts are specified so that a device leaves a channel after hearing no beacons for a pre-specified amount of time.

G.1.3 Incumbent Detection Recovery

Once devices determine that an incumbent has appeared on an in-band channel, they enter the recovery mode of operation. During this mode the device executes the Incumbent Detection Recovery Protocol (IDRP) [4], which allows the network to restore its normal operation in a timely fashion with minimal performance degradation. IDRP offers the network a way to maintain the QoS at an acceptable level while protecting incumbent services. This is particularly important for the RC, as the network operation depends on it.

To address that, one of the key concepts in IDRP is the use of BCs, which allow C-MAC to quickly re-establish communication in the event of an incumbent appearance. BCs are kept in a priority list and are used by a device during the recovery procedure. This way, the recovery procedure can be made very efficient, as all devices on the same channel share the same BC list.

The BC list is built and maintained through out-of-band measurements, which are performed by a device whenever it is not part of any communication. Results of the out-of-band measurements, such as which channels are (un)occupied by incumbents, are communicated through the beacon frames transmitted during the BP. Through this mechanism, devices can learn about the spectrum occupancy in their vicinity. In case traffic load is high and devices do not have a chance to perform out-of-band measurements within a pre-defined amount of time, QPs can be used. Alternatively, a device may deliberately announce in its beacon that it will perform out-of-band measurements for a certain amount of time.

G.2 Self-Coexistence

Self-coexistence mechanisms deal primarily with etiquettes for spectrum sharing for devices operating on the same channel. Given that C-MAC incorporates the use of extended beacon groups as explained in Section IV.D, self-coexistence issues are substantially minimized. Beacons and data transmissions of reachable devices (even if they are not direct neighbors) are all coordinated, and cause negligible, if any, performance degradation.

Mobility (and radio propagation changes), on the other hand, is the detrimental factor. With mobility, devices may come within and out of radio range of each other rapidly, thus disrupting communication through packet collisions due to overlapping superframes.

To address that, C-MAC incorporates a BP merging procedure wherein BPs of devices can be dynamically merged and synchronized, and works as follows. A correctly received beacon that indicates a BPST that is not aligned with a device's own BPST is referred to as an *alien beacon*. The BP defined by the BPST and BP length in an alien beacon is referred to as an *alien BP*. Two BP merging procedures are defined depending upon whether the BPST of the receiving device falls within an alien BP, termed as overlapping BPs, or not, termed as non-overlapping BPs (note that, across channels, C-MAC is able to maintain non-overlapping BPs).

In the overlapping BP case, the device changes its BPST to the BPST of the alien BP. It then adjusts its beacon slot number such that its new beacon slot number is its old beacon slot number plus one, plus the number of the highest occupied beacon slot indicated in any beacon received in the alien BP. After that, the device stops sending further beacons in its previous BP.

In the non-overlapping BP case, the device follows a synchronization rule to dictate the merging and that has been shown to provide very quick convergence [4][7]. In simple terms, this rule states that a device must relocates its beacon to the alien BP only if the alien BPST falls within the first half of the superframe. Obviously, a device does not relocate to the alien BP if a beacon received in that alien BP indicates that the other devices have decided to merge first.

The extended beacon group concept together with the BP merging procedure allows C-MAC to be highly effective in dealing with self-coexistence issues. Extensive simulation results of the synchronization algorithm used in this scheme can be found in [4].

V. PERFORMANCE EVALUATION

We have embarked on analytical, simulations and real prototypical implementation efforts in order to extract key performance attributes of C-MAC. In this section we describe this effort and some preliminary results.

A. Analytical

To understand the maximum throughput that C-MAC can provide, we consider a single transmitter-receiver pair per channel and assume that the transmitter has infinite frames to send to the receiver. We assume the BP length to be fixed at 20ms out of a 100ms superframe size. Similar figures as in 802.11a are adopted for this analysis, and so we further assume that the physical data rate per channel is 54 Mbps (as in 802.11a) and also consider a similar physical layer overhead (e.g., preamble, SIFS time, etc.). Figure 6 depicts the maximum MAC throughput achievable (i.e., throughput available above the MAC) for different frame sizes when 1, 3, and 5 channels are available. As expected, larger frame sizes yield better throughputs because of small overheads. More importantly, we can see that the maximum throughput increases significantly as more channels are added to the network, hence enabling C-MAC to efficiently use the available spectrum.

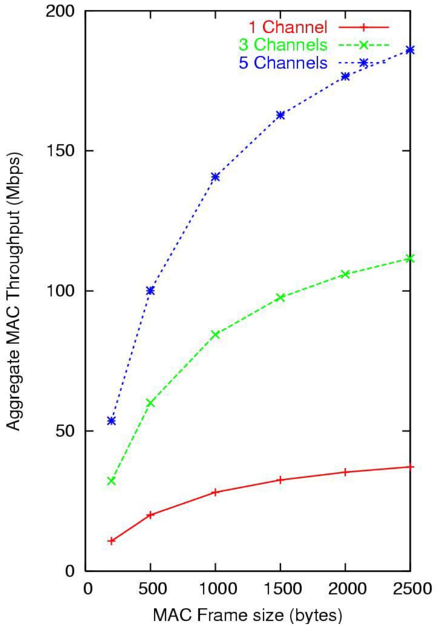

Figure 6 – Maximum C-MAC aggregate throughput above the MAC layer

B. Simulations

In this section, we describe results of simulations conducted to characterize the inter-channel synchronization properties of C-MAC as described in Section IV.E. By simulating convergence times for random topologies and random start times, we plan to ascertain how robust the beaconing mechanism is even under extreme conditions. One such extreme condition is when two or more otherwise uncoordinated C-MAC networks come within radio range of each other.

The simulation consists of randomly placing variable number of nodes in a space of size 50m by 50m. Each node has a communication range of 25m. As discussed in Section IV, initially each node has a random BPST. After the beacons in each channel converge to one BP, one node from each of these channels (a proxy) visits the RC. Within the RC, each of these proxies initially sends beacons at random times, which then align into one BP. In addition, for a given device density (that is number of nodes per channel, on x-axis), 1000 random topologies are simulated and the average convergence time is plotted.

As can be seen from the results, beaconing convergence is very stable and converges within a few superframes, on average.

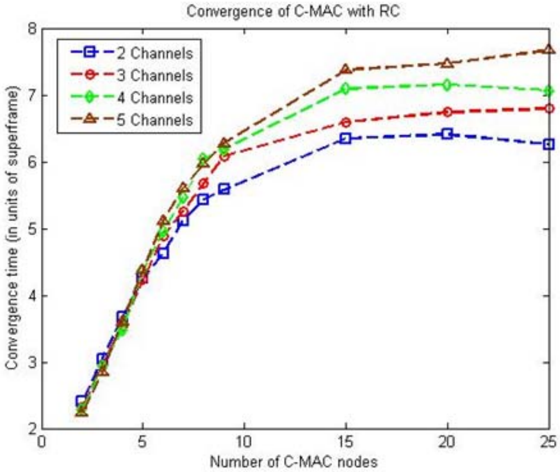

Figure 7 - Simulation results showing convergence time of C-MAC BPs with dynamic RC, with respect to number of nodes per channel (x-axis) and number of channels

C. Prototyping

We have been developing a comprehensive CR prototype implementation in hardware, which includes spectrum sensing, frequency agility, multi-channel operation and so on. For that, we have selected the UNII frequency bands and have utilized IEEE 802.11a cards and Linksys WPC55AG adapters built with Atheros chipsets. In our testbed, the incumbent signal is generated by Rohde&Schwarz signal generator that can transmit a sine-wave signal with controllable center frequency on any of the 13 channels considered. We also use the MadWiFi Linux device driver that is supported by the Atheros chipset. Figure 8 shows a partial view of this setup which, in this case, is comprised of 4 Dell (named node A, B, C, and D) laptops equipped with the aforementioned hardware. In this Figure, the incumbent signal generator is the device in between the four laptops.

We have implemented a superset of C-MAC in our prototype, including the superframe structure with the BP, the inter-channel coordination mechanism, the notion of RC, and a number of coexistence mechanisms (e.g., IDRP). Figure 9 highlights the display of the laptop representing node D, which depicts 3 selected channels and where channel 44 is the RC. Through the RC, node D can quickly discover other nodes and this is shown in Figure 9. Since this prototype and its results shall be described in greater detail during DySPAN 2007, for the purpose of this paper we do not describe it further. We note, however, that more information on the spectrum sensing component of this prototype can be found in [27].

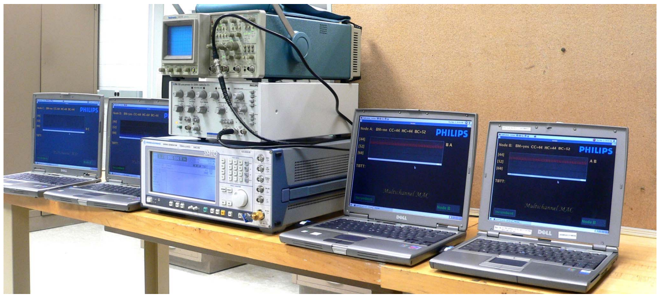

Figure 8 – Prototypical CR testbed including C-MAC

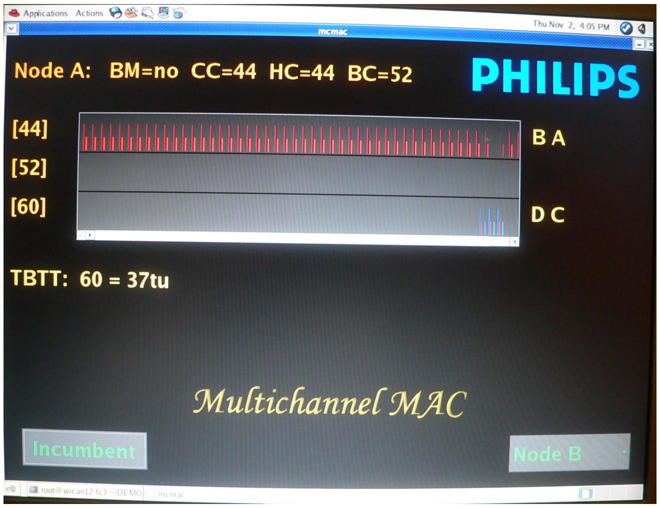

Node A: BM=no CC=44 HC=44 BC=52

TBTT: 60 = 37tu

Figure 9 – Snapshot of node D's display illustrating data transmission across channels

VI. CONCLUSIONS

In this paper we have introduced the cognitive MAC (C-MAC) protocol for distributed multi-channel wireless networks. C-MAC operates over multiple channels, and hence is able to effectively deal with, among other things, the dynamics of resource availability due to primary users. A key concept in C-MAC is the use of a dynamic and totally distributed RC which, among other things, is used to support network-wide multicast and broadcast. C-MAC has been evaluated analytically, through simulations and implemented in a real hardware prototype. Performance results are very promising and some of them have been presented here. We believe that C-MAC is novel in many respects and opens up new research directions in MAC protocols for cognitive multi-channel wireless networks.

REFERENCES

[1] J. Mitola et al., "Cognitive Radios: Making Software Radios more Personal," *IEEE Personal Communications*, vol. 6, no. 4, Aug. 1999.

[2] J. Mitola, "Cognitive radio: An integrated agent architecture for software defined radio," *PhD Dissertation*, Royal Inst. Technol. (KTH), Stockholm, Sweden, 2000.

[3] S. Haykin, "Cognitive Radio: Brain-Empowered Wireless Communications," in *IEEE JSAC*, vol. 23, no. 2, Feb. 2005.

[4] C. Cordeiro, K. Challapali, and M. Ghosh, "Cognitive PHY and MAC Layer for Dynamic Spectrum Access and Sharing of TV Bands," in *IEEE International Workshop on Technology and Policy for Accessing Spectrum* (TAPAS), August 2006

[5] K. Challapali, C. Cordeiro, and D. Birru, "Evolution of Spectrum-Agile Cognitive Radios: First Wireless Internet Standard and Beyond," in *2^{nd} ACM International W*ireless Internet Conference (WICON), August 2006.

[6] M. Marsan and D. Roffinella, "Multichannel Local Area Network Protocols," *IEEE JSAC*, Vol. 1, No. 5, November 1983.

[7] IEEE 802.22 draft standard, "IEEE P802.22^{TM}/D0.1 Draft Standard for Wireless Regional Area Networks," http://www.ieee802.org/22/, doc. no. 22-06-0068-00-0000, May 2006.

[8] S. Seidel and R. Breinig, "Autonomous Dynamic Spectrum Access System Behavior and Performance," in *IEEE DySPAN*, Nov. 2005.

[9] L. Ma, X. Han, and C.-C. Shen, "Dynamic Open Spectrum Sharing for Wireless Ad Hoc Networks," in *IEEE DySPAN*, Nov. 2005.

[10] Q. Zhao, L. Tong, and A. Swami, "Decentralized Cognitive MAC for Dynamic Spectrum Access," in *IEEE DySPAN*, Nov. 2005.

[11] Z. Tang and J. J. Garcia-Luna-Aceves, "Hop-Reservation Multiple Access (HRMA) for Ad Hoc Networks," in *IEEE Infocom*, 1999.

[12] A. Tzamaloukas and J. J. Garcia-Luna-Aceves, "A Receiver-Initiated Collision-Avoidance Protocol for Multi-Channel Networks," in *IEEE Infocom*, 2001.

[13] J. Li, Z. Haas, M. Sheng, and Y. Chen, "Performance Evaluation of Modified IEEE 802.11 MAC for Multi-Channel Multi-Hop Ad Hoc Network," *International Conference on Advanced Information Networking and Applications (AINA)*, March 2003.

[14] S. Chaudhuri, R. Kumar, and A. Saha, "A MAC Protocol for Multi Frequency Physical Layer," *Technical Report*, Rice University, Texas, Jan. 2003.

[15] P. Porwal and M. Papadopouli, "On-Demand Channel Switching for Multi-Channel Wireless MAC Protocols," *Technical Report TR04-024*, University of North Carolina at Chapel Hill, 2004.

[16] J. So and N. Vaidya, "Multi-Channel MAC for Ad Hoc Networks: Handling Multi-Channel Hidden Terminals Using a Single Transceiver," in *ACM Mobihoc*, May 2004.

[17] W.-C. Hung, K. Law, A. Leon-Garcia, "A Dynamic Multi-Channel MAC for Ad-Hoc LAN," in *21st Biennial Symposium on Communications*, April 2002.

[18] A. Nasipuri and S. Das, "Multichannel CSMA with Signal Power-based Channel Selection for Multihop Wireless Networks," in *IEEE VTC*, Sept. 2000.

[19] N. Jain and S. Das, "A Multichannel CSMA MAC Protocol with Receiver-based Channel Selection for Multihop Wireless Networks," in *IEEE ICCCN*, Oct. 2001.

[20] S.-L. Wu, C.-Y. Lin, Y.-C. Tseng, and J.-P. Sheu, "A New Multi-Channel MAC Protocol with On-Demand Channel Assignment for Multi-Hop Mobile Ad Hoc Networks," in *I-SPAN*, 2000.

[21] A. Nasipuri, J. Zhuang, and S. Das, "A Multichannel CSMA MAC Protocol for Multihop Wireless Networks," in *IEEE WCNC*, 1999.

[22] J. Chen, S. Sheu, and C. Yang, "A new multichannel access protocol for ieee 802.11 ad hoc wireless lans," in *IEEE PIMRC*, pp. 2291 – 2296, 2003.

[23] A. Tzamaloukas and J. Garcia-Luna-Aceves, "Channel-hopping multiple access," in IEEE ICC, June 2000.

[24] P. Bahl, R. Chandra, and J. Dunagan, "SSCH: Slotted seeded channel hopping for capacity improvement in IEEE 802.11 ad-hoc wireless networks," in *ACM MobiCom*, Sept. 2004.

[25] D. Raychaudhuri and X. Jing, "A spectrum etiquette protocol for efficient coordination of radio devices in unlicensed bands," in *14th IEEE PIMRC*, pp. 172–176, Sept. 2003.

[26] C. Cordeiro, H. Gossain and D. Agrawal, "Multicast over Wireless Mobile Ad Hoc Networks: Present and Future Directions", in *IEEE Network*, SI on Multicasting: An Enabling Technology, January/February 2003.

[27] H. Kim, C. Cordeiro, K. Challapali, K. Shin, "An Experimental Approach to Spectrum Sensing in Cognitive Radio Networks with Off-the-Shelf IEEE 802.11 Devices," in *IEEE Workshop on Cognitive Radio Networks*, in conjunction with IEEE CCNC, Jan. 2007.

[28] J. Wang, Y. Fang, and D. Wu, "A power-saving multi-radio multi-channel MAC protocol for wireless local area networks," in *IEEE Infocom*, Barcelona, April 2006.