

# A Bandwidth Sharing Approach to Improve Licensed Spectrum Utilization



PANAGIOTIS PAPADIMITRATOS, SRIVATSAN SANKARANARAYANAN, AND  
AMITABH MISHRA, VIRGINIA TECH

## Abstract

The spectrum of deployed wireless cellular communication systems is found to be underutilized, even though licensed spectrum is at a premium. To efficiently utilize the bandwidth left unused in a cellular system, the primary system (PRI), we propose an overlaid ad hoc secondary network (ASN) architecture, with the ASN operating over the resources left unutilized by the PRI. Our basic design principle is that the ASN operates in a nonintrusive manner and does not interact with the PRI. In this article we present the ad hoc secondary medium access control (AS-MAC) protocol to enable PRI-SEC interoperation, address a number of technical challenges pertinent to this networking environment, and evaluate the performance of the AS-MAC. In a single-hop ASN the AS-MAC transparently utilizes 75 percent of the bandwidth left unused by the PRI, while in multihop ASNs, due to spatial reuse, the AS-MAC can utilize up to 132 percent of the idle PRI resources in our experiments.

## Introduction

There is a strong belief that the spectrum in both the public and private sectors in the United States is getting scarce. However, recent measurements for cellular systems in major metropolitan areas [1] suggest that in several frequency bands spectrum utilization is very low for extended periods of time. In other words, spectrum is available but inefficiently utilized; thus, new modulation and coding techniques alone cannot address the problem at hand. Spectrum sharing or pooling, first introduced in [2], is a promising approach that can enable access to the same spectrum by two disjoint systems. In our context, the owner of the spectrum, which we term the primary system (PRI), can allow a secondary system (SEC) to utilize the portion of the same spectrum left unused by the PRI. The fundamental constraints for such an architecture are that the SEC operates strictly over the resources (i.e., bandwidth) left unutilized by the PRI, SEC operation does not degrade the performance of the PRI, and there is no signaling exchanged between the PRI and SEC. Overall, SEC operation must be transparent to the PRI.

Based on these requirements, we propose an architecture with a multihop ad hoc network, the *ad hoc secondary network* (ASN), as the SEC system, overlaid on a time-/frequency-domain multiple access (TDMA/FDMA)-based Global System for Mobile Communications (GSM) cellular network. Figure 1 illustrates the PRI-SEC system: for the PRI, mobile stations (MSs) communicate with the base station (BS), while for the ASN, ad hoc nodes (ANs) form a multihop peer-to-peer topology. ANs first determine the communication structure of the PRI, identify the available resources (i.e., time slots within each of the cell's frequency bands), and utilize only this bandwidth to communicate. One example of a deployment scenario is a cellular provider that leases its unused spectrum

during low cellular traffic periods (e.g., night). The SEC can be, for example, a mesh network providing peer-to-peer communication or wireless Internet access.

In the literature [3] considers a PRI-SEC system, proposing a HIPERLAN SEC based on orthogonal frequency-division multiplexing (OFDM) modulation operating along with a PRI GSM. Our work is significantly different, as we develop an ad hoc SEC system that operates without fixed infrastructure, and address a number of technical issues related to interoperation with the PRI GSM, such as SEC traffic transmission. Moreover, our design has the advantage of being largely independent of the physical layer. In [4] a medium access control (MAC) protocol for a single-channel PRI-SEC configuration is designed under the assumption the system can predict "spectrum holes" used to transfer packets. Our work is not dependent on the prediction of resource availability, thus ensuring noninterference between the PRI and SEC, while being capable of operating over multiple channels.

Beyond the PRI-SEC context, a number of multichannel MAC (MMAC) protocols were proposed. However, those are either inapplicable or inefficient and thus impractical in the PRI-SEC setting: [5, 6] require nodes equipped with a number of transceivers equal to the number of channels, and [7] requires three transceivers. A solution with two transceivers, one of them tuned on the control channel, was proposed in [8]. This scheme introduces a new control packet called RES, which contains channel reservation information similar to clear-to-send (CTS) and helps avoid collisions. Our system operates with a single transceiver, along with a channel sensing module, and reduces the control overhead, as the RES control packet is not beneficial [9].

## System Architecture and Overview

A set,  $C$ , of channel pairs (frequency bands) is allocated to each cell of the PRI, out of  $C_{total} = 124$  available GSM bands [10]. Each uplink (MS to BS) or downlink (BS to MS) channel is divided into frames, with  $T_f$  time slots/frame. Time slots are mapped onto logical *control* and *traffic* channels, with voice or data communication taking place on traffic channels, and signaling information transmitted on the control channels. The frequency correction channel (FCCH) and synchronization channel (SCH) enable the MSs to achieve time synchronization with the BS. The physical frequency that carries FCCH is denoted the beacon frequency (or channel). The broadcast control channel (BCCH) provides the location area identity (LAI), cell identity (CI), and cell channel description (CCD). The combination of LAI and CI, termed the cell global identity (CGI), uniquely identifies a cell, while the CCD provides a list of the cell frequencies (i.e.,  $C$ ). The common signaling channels are not encrypted; not only MSs but also ANs within the cell radius can readily receive the BS signaling.

The first challenge for ANs is to determine the slot boundaries: they can obtain this information by decoding the FCCH

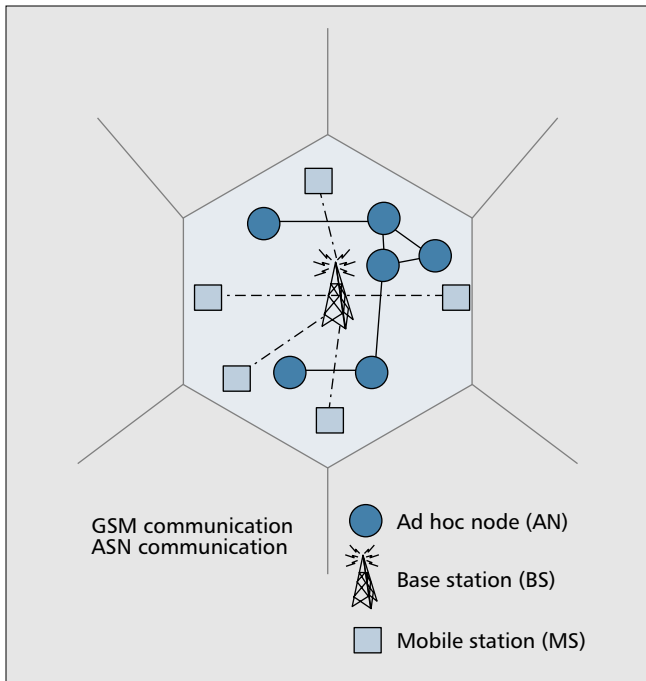


FIGURE 1. Proposed PRI-SEC architecture.

signaling from the BS, a task that should be performed periodically; for example, MSs obtain timing information from the BS twice per second (when a call is in progress). Then ANs decode SCH and BCCH, with the latter providing the CGI and CCD of the current cell. Based on CGI, ANs determine if they are allowed to share the cell PRI resources with CCD providing the frequencies of the cell. The BCCH and SCH signaling enables ANs to obtain the slot boundaries.

Once  $C$  and the slot timing are determined, ANs create and maintain an up-to-date map of available time slots on the downlinks. The system presented in this article utilizes resources available on the downlinks only. This is so because determining the slot boundaries of the uplinks would require cooperation from the PRI (i.e., the BS). We assume that ANs are equipped with a *sensing module* capable of wideband spectrum sensing [11]. For our system, it suffices that the sensing module detects the presence of a signal (i.e., energy level above a threshold) within each of the  $C$  bands. For a short period of time at the beginning of each slot, the sensing module detects if a PRI transmission takes place. If no transmission is sensed, it is definite that the PRI GSM will not utilize the current slot. ASN can then utilize it without the risk of colliding with PRI traffic. We emphasize that without per-slot sensing, use of only unused PRI resources by the ASN cannot be ensured. This is because prediction of future PRI usage cannot, in principle, be flawless.

AS-MAC operates above GSM MAC, as shown in Fig. 2a, to isolate the ASN network layer from GSM. The ANs have only one transceiver, while, in general, multiple GSM channels are available. AS-MAC provides for the selection of one among those channels. To do so, a handshake is necessary between the sending and receiving ANs: the sender provides candidate channels, and the receiver selects a desirable channel. This exchange of information is performed across a commonly agreed (within the ASN) channel, which we denote the control channel (CC). The actual data transmission takes place across the selected channel among the remaining ones in  $C$ , which we denote the data channel (DC).

Once the DC is selected, AS-MAC transmits the data. In general, ASN packets can be larger than the number of bits that can be transmitted within an interval of PRI inactivity; thus, data may need to be fragmented. The challenge here is that transmission must take place in available slots that in general are not consecutive. Since the occupancy (availability) of slots depends on PRI traffic, time progress of the ASN protocol, in our case AS-MAC, must take place only when PRI slots are free. When there is PRI activity, the state of the ASN protocol must essentially freeze. Consider, for example, Fig. 2b, showing a single time-slotted GSM downlink consisting of eight slots numbered 0 to 7; slots 0, 3, and 7 used by the PRI, and slots 1, 2, 4, 5, and 6 left unutilized. If a message transmission is to occupy three slots, starting from slot 1, counting those slots must “stop” during slot 3 and resume with slot 4.

### AS-MAC Protocol Operation

Among the set,  $C$ , of GSM cell frequencies, ANs select the one with the highest index, other than the beacon frequency, as the ASN CC. The set  $C_d$  of remaining channels are the DCs. Through sensing, ANs update locally information on PRI slot usage over the last  $u$  slots so that they dynamically select the preferred data channel for each packet transmission. The sensing module is also utilized for sensing ASN transmissions, after the primary signal sensing, for a period of  $\tau_{SEC}$  s. We denote this *secondary sensing*, performed on the ASN CC and DCs so that information on the DCs currently in use by other ANs can be maintained.

With the resource availability information at hand, AS-MAC provides the means for nodes to first agree on a CD, through a handshake that involves the exchange of three control messages, request-to-send (RTS), clear-to-send (CTS), and reservation (RES), transmitted in this order, as explained below. We found in [9] that the RES may not be necessary, thanks to the sensing module. We identified two versions of AS-MAC, one that uses RES and we denote AS-MAC<sub>1</sub>, and one without RES denoted AS-MAC<sub>2</sub>. Since the latter is found to be more efficient, we discuss this variant below, referring to AS-MAC<sub>1</sub> and AS-MAC<sub>2</sub> interchangeably unless otherwise noted.

Figure 3a illustrates the AS-MAC operation: the sender AN waits for a free slot on CC and transmits (unicasts) an




FIGURE 2. a) ASN protocol stack; b) GSM slot utilization by ASN nodes.

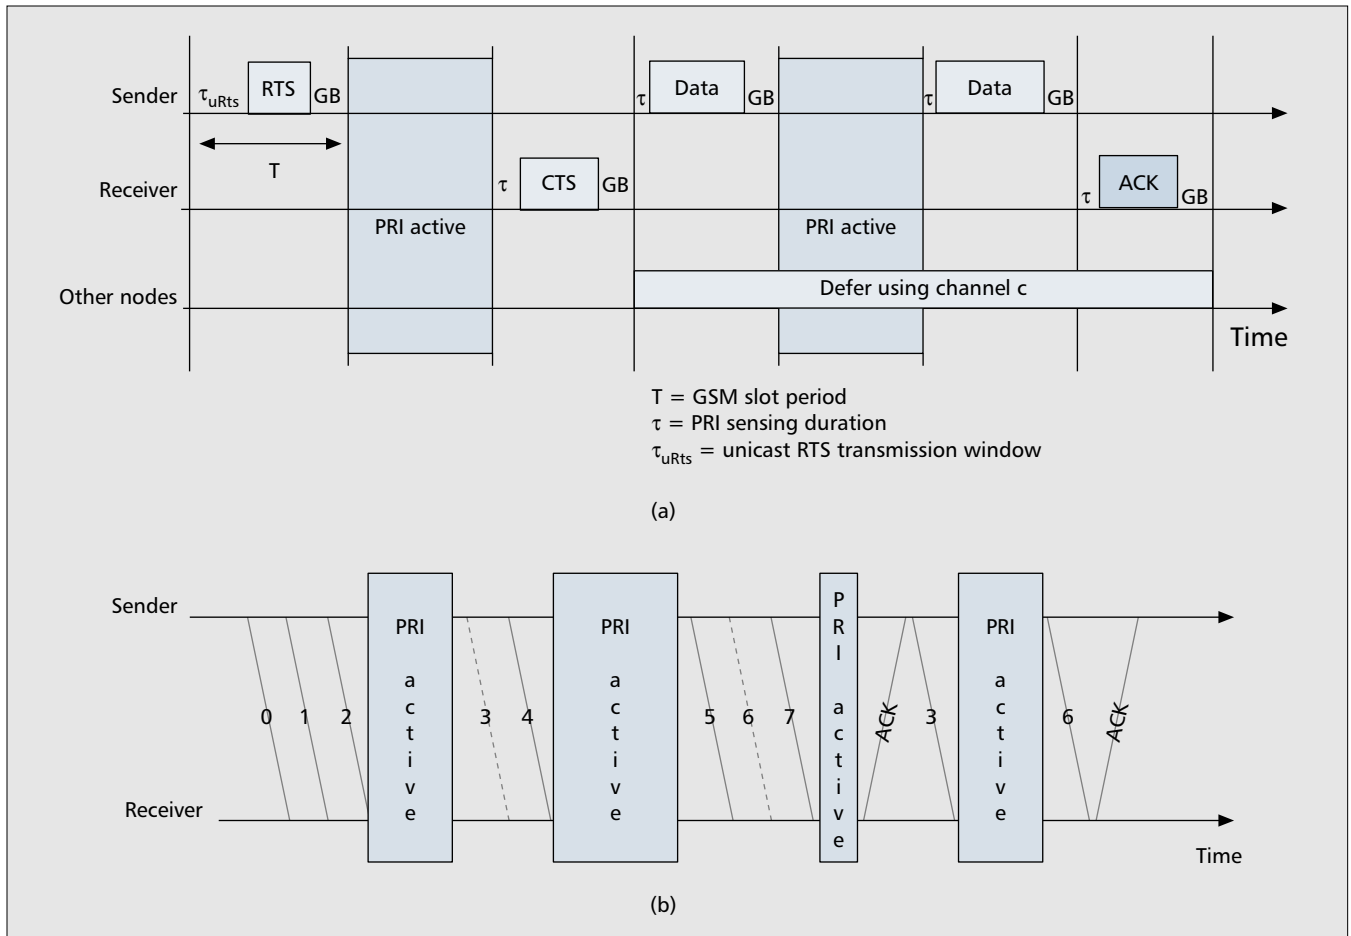


FIGURE 3. a) AS-MAC transmission; b) AS-MAC error recovery.

RTS for an intended receiver. The RTS transmission takes place at time  $\tau_{uRTS}$  randomly selected from a uniform distribution  $U(W_1, W_1 + W_{uRTS})$ , after the slot beginning, so that collisions among RTS packets are reduced. The RTS contains a bit map of the channels in  $C_d$  preferred by the sender, and the number of PRI slots needed to transmit the packet, which we denote NAV. The idea behind this NAV is similar to that in 802.11, except that now it is specified in terms of the number of free slots. If CTS is not received, RTS is reattempted a number of times. If successful, the CTS provides the receiver and sender ID, NAV, and channel for communication,  $c$ . The same information is provided by the RES, if sent. The selected channel is the one with the maximum number of free slots, based on PRI traffic among those preferred by both the sender and receiver. A channel is selected only if it is either sensed free or known to be free based on prior receipt of CTS (or RES) packets from other nodes. Nodes receiving CTS (or RES) can defer sending an RTS to a busy node involved in the corresponding data transfer, but, more important, they can unambiguously “prohibit” themselves from using  $c$  until at least NAV free PRI slots elapse on  $c$ . We emphasize that the sender does **not** reserve a channel for a fixed duration of time, as is the case with 802.11 and MMAC protocols. In the PRI-SEC context, the AN cannot know the future channel/slot usage of the primary system, and thus cannot predict the transmission duration. Rather, AS-MAC sets the number of free slots required to transmit the data packet, and ANs decrement their counters only when a slot on  $c$  is PRI-free.

After the RTS/CTS (RES) handshake is completed, the sender fragments the packet and transmits the fragments, identified by an increasing sequence number, successively on free slots of the data channel. Figure 3a illustrates that AS-MAC

freezes its state, remaining idle in all slots utilized by the PRI. An acknowledgment (ACK) is expected from the sender upon transmission of all fragments. The sender decrements a *count* field in each fragment header to provide robustness to network impairments. The receipt of a single fragment, not necessarily the first transmitted one, suffices for the receiver to “schedule” the ACK transmission, independent of how many packet fragments are actually received. If all fragments are received, the packet is reconstructed and passed to the upper layer, and the receiver acknowledges the receipt. Otherwise, the ACK contains the IDs of the lost fragments<sup>1</sup> so that the sender can retransmit only those. This way, network overhead can be reduced compared to full retransmissions, an important aspect in resource constrained environments. Figure 3b illustrates an example of error recovery: a packet consists of eight fragments, with IDs from 0 to 8; fragments 3 and 6 (shown in dotted lines) are lost; the first ACK notifies the sender, which then retransmits them to complete the packet transfer. A protocol-selectable maximum number of retransmissions are attempted, with the packet transmission aborted if they all fail or when no PRI-free slots are available.

## Performance Evaluation

We evaluate our system, studying the improvement in spectrum utilization due to the ASN within a single cell of the PRI GSM system. We experiment with three ASN topologies, in all cases fully residing within the GSM cell:

<sup>1</sup> The AS-MAC maximum transmission unit (MTU) and ASN physical layer determine the maximum size of the ACK.

- A fully connected network of 20 ANs
- A  $10 \times 10$  grid topology of 100 ANs
- A random topology with 100 ANs having uniformly random locations within a  $1000 \text{ m} \times 1000 \text{ m}$  square area

In the experiments presented here ANs are static (no mobility). The ANs' nominal transmission and sensing ranges are 250 m and 625 m, respectively, with a 10 dB signal-ti-interference-plus-noise ratio (SINR) threshold for successful reception. ANs are equipped with one transceiver and a sensing module. Moreover, for this performance evaluation, we consider the variant of AS-MAC without the RES message (AS-MAC<sub>2</sub>). ANs transmit each packet to a randomly selected neighbor. We evaluate ASN in saturation conditions, with each AN always having a packet with payload of 280 bytes to transmit; as a result, we ensure that the ASN has sufficient load to fully utilize the available resources.

There are  $C = 8$  PRI GSM channels in use within the cell, one of them fixed for the ASN control traffic and the remaining seven channels for data traffic. Each PRI voice call occupies one time slot of every frame on the allocated uplink and the downlink. Calls arrive according to a Poisson process with aggregate rate  $\gamma$ , and call holding times are exponentially distributed with mean  $1/\mu$ . Typical values of  $1/\mu$  are on the order of minutes [10]. Each new call is allocated randomly to an available slot. We vary the call arrival rate to generate GSM loads, or percentages of utilized slots, across all PRI channels, which we denote PRI utilization or  $PRI_U$ , with average values of 12.5, 25, 50, 75, and 87.5 percent over the simulation duration (600 s).

We evaluate the following metrics:

- Available bandwidth utilization (BU), the fraction of the number of slots, across all ASN data channels, utilized by ANs, over the number of slots left unused by the PRI
- Spectrum utilization improvement (SUI) =  $(PRI\_ASN_U - PRI_U)/PRI_U$ , where  $PRI\_ASN_U$  is the percent of slots utilized when both the PRI and ASN are deployed

BU quantifies the efficiency of AS-MAC in utilizing the available resources, while SUI provides an overall picture of utilization effectiveness. Note that both metrics are independent of the ASN physical layer (PHY) and thus can serve as a benchmark for any ASN PHY implementation. We assume for illustration purposes the ASN PHY data rate to be equal to the GSM PHY rate. Clearly, the higher the ASN PHY data rate, the higher the throughput will be, as PHY protocols other than the GSM one could be used by the ASN. All data points, shown with 95 percent confidence intervals, are the averages over at least five randomly seeded runs.

Figures 4a–b show the BU and SUI metrics as a function of  $PRI_U$  for the single-hop and random multihop topologies, respectively. In all three cases the ANs operate in saturation conditions. In single-hop ANs (Fig. 4a) BU is up to 75 percent, or, in other words, *three out of four* available PRI slots are utilized for ASN data communication. Then in Fig. 4b BU is up to 132 percent for random multihop ANs. BU exceeds 100 percent, as in multihop topologies a single idle PRI slot can be utilized in principle by two or more simultaneous ASN transmissions. The degree of spatial reuse of the available bandwidth depends on the ANs' transmission and sensing range, as well as the multihop topology characteristics.

In both figures, BU remains practically constant as  $PRI_U$  increases, showing that AS-MAC remains efficient when either a significant or small portion of the PRI resources are available. SUI in Fig. 4a–b decreases as  $PRI_U$  increases, with values ranging from 250.3 to 5.37 percent for the single-hop ASN, and from 418 percent to 10.5 percent for the random multihop topology. Essentially, SUI shows the improvement from deploying an ASN as a function of the PRI load. As  $PRI_U$  increases, the resources available for ASN communication decrease; thus, the additional utilization offered/achieved by the ASN is unavoidably bounded from above by  $100 - PRI_U$ .

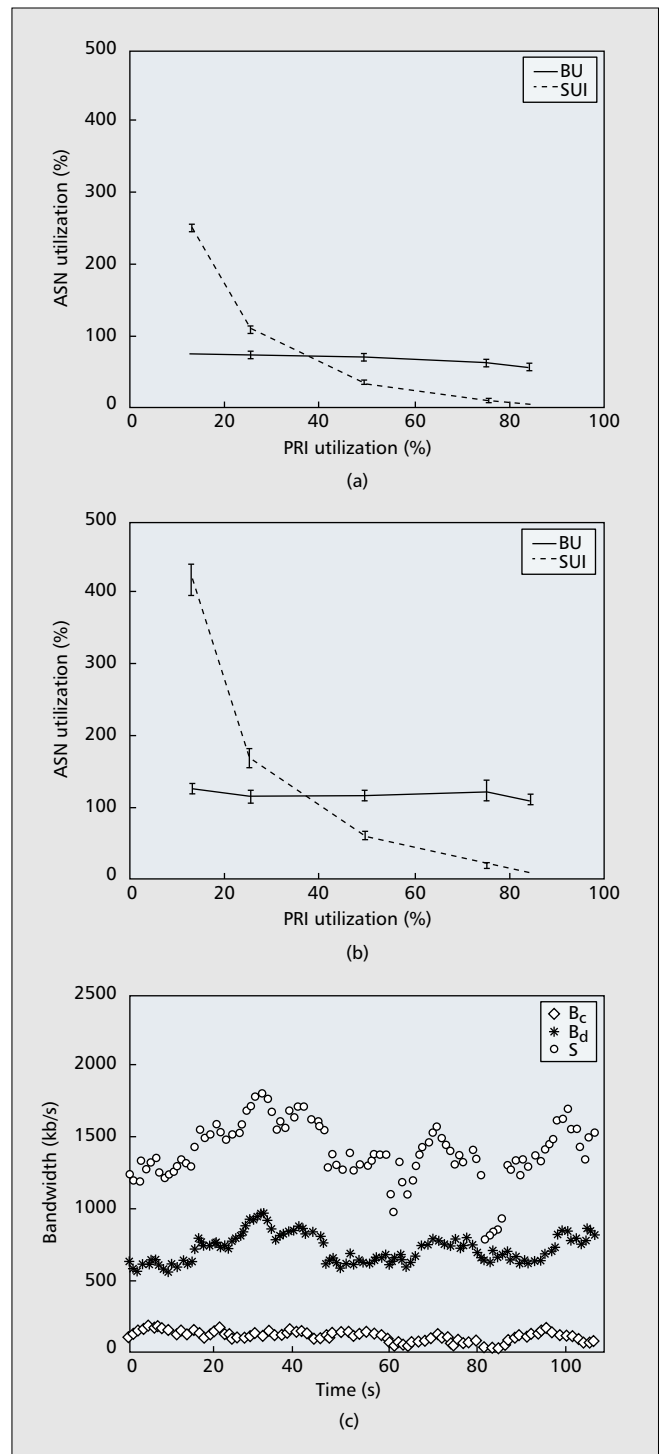


FIGURE 4. a) 20-node single-hop ASN: utilization performance as a function of PRI utilization (load); b) 100-node random multihop ASN: utilization performance as a function of PRI utilization (load); c) ASN throughput trace for a  $10 \times 10$  grid.

Nevertheless, when  $PRI_U$  is low, AS-MAC realizes the benefits of bandwidth sharing.

## Discussion and Future Work

Figure 4c traces the variation of the ASN throughput,  $S$ , over time for the  $10 \times 10$  grid when the PRI offered load is set to 50 percent, with  $B_c$  and  $B_d$  the available bandwidth on the control (CC) and data channels ( $C_d$ ), respectively, all in kilo-



bits per second.  $B_c$  and  $B_d$  can be directly calculated based on the number of free slots on the CC and DCs, respectively. Due to the time varying nature of the PRI traffic,  $B_c$  and  $B_d$  fluctuate. At  $t = 30$  s, there is a sharp increase in  $S$  because of a corresponding increase in  $B_d$ . A similar peak occurs at  $t = 38$  s. At  $t = 62$  s,  $B_c$  is very low, leading to a low  $S$  as well, despite the relatively high  $B_d$ . The same phenomenon occurs at  $t = 83$  s. It is thus evident that AS-MAC performance is dependent on how  $B_c$  and  $B_d$  vary over time. Due to the dynamic nature of PRI traffic, low  $B_c$  degrades  $S$ . A solution to this problem is for ANs to dynamically switch to a new CC based on the observed PRI load. This is possible exactly because the sensing module provides ANs with PRI activity information. We are currently investigating the design of robust algorithms for selecting an alternate ASN CC.

Another important aspect is to ensure that SEC operation does not degrade PRI performance. Sensing of idle periods alone, however, does not suffice, because the timing information ANs obtain from the BS signaling can be offset due to the signal propagation delay from the BS to the AN. Unlike MSs, which overcome such timing offsets with the assistance of the BS, ANs cannot interact with the PRI to correct their synchronization information. The selection of an appropriate guard band,  $GB$ , bounds the amount of additional interference imposed by all ANs to adjacent time slot BS-to-MS transmissions. As a result, AN transmissions are guaranteed not to interfere with PRI transmissions.  $GB$  along with  $\tau$ , the PRI sensing period, which we assume to be less than 15  $\mu$ s (e.g., the currently available 802.11 DSSS PHY layer sensing period), leave 90 percent or more of an idle PRI slot available for ASN traffic.

As for issues related to ASN transceivers, the turnaround time (i.e., the period of time needed for a transceiver to switch from transmitting to receiving mode and vice versa) can easily be accommodated. In our system, such transitions need to occur at the PRI slot boundaries, with the aggregate of the GSM guard band (15  $\mu$ s), the PRI sensing period, and  $GB$  well above, for example, the 5  $\mu$ s and 10  $\mu$ s receive-to-transmit and transmit-to-receive turnaround times for 802.11g DHSS. Another concern is the time needed to dynamically switch a transceiver to different channels at different points in time. This is necessary in AS-MAC after an RTS-CTS handshake and after the transmission of a packet, when the sender and receiver switch to the CC, with these periods on the order of 60–70  $\mu$ s or less. However, the channel switching time allowed in 802.11 is 224  $\mu$ s, roughly three times higher. To allow ample time for the channel switch, the sender and receiver freeze their operation in the next slot, if it is free, and resume the protocol operation thereafter.

As ANs can potentially receive multiple FCCH messages from different BSs, they can choose a beacon frequency and resources associated with it. To extend the operation of the ASN and enable intercell AN communication, a subset of the ANs will have to tune to multiple cells and act as gateways. We are investigating the operation of the ASN in such a setting. Finally, we note that versatile interoperation of ASN and PRI necessitates that ANs know if the PRI allows sharing its unutilized resources within the underlying cell. In fact, distinguishing among multiple providers would also be necessary. We advocate that the PRI provide such information to the SEC offline, in the form of a service agreement. A “map” of cell identities (e.g., CGIs) that allow resource sharing can be stored in each AN.

## Conclusions

This article is concerned with the efficient utilization of the resources (spectrum) of deployed wireless communication systems, based on a bandwidth sharing approach. The basic

idea is to deploy a secondary system that dynamically and transparently operates only over the resources left unutilized by the licensed user of the spectrum, the primary system. We propose the AS-MAC protocol as the focal point of a PRI-SEC architecture, with a GSM cellular system as the PRI and an ad hoc network as the SEC. Our AS-MAC is shown to efficiently utilize the available resources, with a utilization factor from 75 to 132 percent, due to spatial reuse in the latter case. In spite of a number of technical issues whose solution we outline here and currently investigate, we believe approaches such as the one proposed here based on AS-MAC can efficiently utilize available spectrum, and lead to practical deployment of secondary wireless peer-to-peer and mesh networks.

## References

- [1] [http://www.fcc.gov/sptf/files/SEWGFinalReport\\_1.pdf](http://www.fcc.gov/sptf/files/SEWGFinalReport_1.pdf), FCC Spectrum Policy Task Force, Nov. 2002.
- [2] J. Mitola, III, “Cognitive Radio for Flexible Mobile Multimedia Communications,” *IEEE Int’l. Wksp. Mobile Multimedia Commun.*, Nov. 1999, pp. 3–10.
- [3] T. Weiss and F. Jondral, “Spectrum Pooling: an Innovative Strategy for the Enhancement of Spectrum Efficiency,” *IEEE Commun. Mag.*, Mar. 2004.
- [4] V. Syrotiuk et al., “Dynamic Spectrum Utilization in Ad Hoc Networks,” *Elsevier Comp. Net. J.*, Dec. 2004.
- [5] A. Nasipuri and R. Das, “Multichannel CSMA with Signal Power-Based Channel Selection for Multihop Wireless Networks,” *IEEE VTC*, Sept. 2000.
- [6] N. Jain, R. Das, and A. Nasipuri, “A Multichannel CSMA MAC Protocol with Receiver-Based Channel Selection for Multihop Wireless Networks,” *10th Int’l. Conf. Comp. Commun. and Net.*, Oct. 2001.
- [7] J. Pathmasuntharam, A. Das, and A. Gupta, “Primary Channel Assignment based MAC (PCAM) — A Multichannel MAC Protocol for Multihop Wireless Networks,” *WCNC 2004*, Mar. 2004.
- [8] S. Wu, C. Lin, Y. Tseng, and Y. Sheu, “A New Multichannel MAC Protocol with On-Demand Channel Assignment for Multihop Mobile Ad Hoc Networks,” *Int’l. Symp. Parallel Architectures, Algorithms and Networks*, Dec. 2000.
- [9] S. Sankaranarayanan, P. Papadimitratos, and A. Mishra, “Enhancing Wireless Spectrum Utilization with a Cellular — Ad Hoc Overlay Architecture,” *IEEE MILCOM*, Atlantic City, NJ, Oct. 17–20, 2005.
- [10] B. Mouley and B. Pautet, “The GSM System for Mobile Communication,” *Cell and Systems*, Palaiseau, France, 1992.
- [11] D. Cabric, M. Mishra, and W. Brodersen, “Implementation Issues in Spectrum Sensing for Cognitive Radios,” *Asilomar Conf. Signals, Sys., and Comp.*, 2004.

## Biographies

PANAGIOTIS PAPADIMITRATOS [M’99] (papadp@pop.vt.edu) earned his Ph.D. in electrical and computer engineering from Cornell University, Ithaca, New York, in 2005, and his B.Sc. in computer engineering and informatics from Patras University, Greece, in 1998. He joined the Electrical and Computer Engineering Department of Virginia Tech, Blacksburg, in 2004, as a research associate. He is the author of more than 20 technical papers, and his research is concerned with networking protocols, network security, ad hoc and sensor networks, and mobile computing.

SRIVATSAN SANKARANARAYANAN (srivatsan@vt.edu) received a B.Tech. degree in electrical engineering from the Indian Institute of Technology, Madras, in 1999, and an M.S. degree in electrical engineering from The George Washington University, Washington, DC, in 2003. Currently, he is a doctoral student at Virginia Tech, Blacksburg. His research interests are in wireless network architectures and protocols.

AMITABH MISHRA [SM] (mishra@vt.edu) is an associate professor in the Bradley Department of Electrical and Computer Engineering at Virginia Tech, which he joined in August 2000. At Virginia Tech he works in the areas of dynamic spectrum management, protocols and systems for heterogeneous networking, and wireless sensor network architecture. Currently he is working on cross-layer design optimization of the sensor networking protocol stack, media access control algorithms for cellular-ad hoc networking, and wireless security. His research has been sponsored by National Security Agency, DARPA, NSF, NASA, and Raytheon. From 1987 to 2000 he was a member of technical staff at Lucent Technologies — Bell Laboratories, Naperville, Illinois, where his focus was on architecture and performance of communication applications running on 5ESS switches. GPRS, CDMA2000, and UMTS were a few of the areas he worked on while with Bell Laboratories. He received his B.Eng. and M.Tech. degrees in electrical engineering from Government Engineering College, Jabalpur, and Indian Institute of Technology, Kharagpur, in 1973 and 1975, respectively. He obtained his M.Eng. and Ph.D. in 1982 and 1985, respectively also in electrical engineering, from McGill University, and an M.S. in computer science in 1996 from the University of Illinois at Urbana-Champaign. He is a member of ACM, SIAM, and ASEE.