


A DECISION-THEORETIC FRAMEWORK FOR OPPORTUNISTIC SPECTRUM ACCESS

QING ZHAO, UNIVERSITY OF CALIFORNIA AT DAVIS
ANANTHRAM SWAMI, U.S. ARMY RESEARCH LABORATORY



The authors identify basic components, fundamental trade-offs, and practical constraints in opportunistic spectrum access. A decision-theoretic framework based on the theory of partially observable Markov decision processes is introduced.

ABSTRACT

Built on a hierarchical access structure with primary and secondary users, opportunistic spectrum access improves spectrum efficiency while maintaining compatibility with legacy wireless systems. The basic idea is to allow secondary users to exploit instantaneous spectrum availability while limiting the interference to primary users. In this article, we identify basic components, fundamental trade-offs, and practical constraints in opportunistic spectrum access. We introduce a decision-theoretic framework based on the theory of partially observable Markov decision processes. This framework allows us to systematically tackle the optimal integrated design and quantitatively characterize the interaction between signal processing for opportunity identification and networking for opportunity exploitation. A discussion of open problems, potential applications, and recent developments is also provided.

INTRODUCTION

Measurements of actual spectrum usage have revealed the pervasiveness of idle frequency bands in the seemingly crowded radio spectrum [1]. Due to bursty arrivals of wireless applications and guard bands in space, much of the prized spectrum lies unused at any given time and location. Shown in Fig. 1 is a wireless LAN traffic measurement, indicating 75 percent idle time during an active FTP session [2]. For voice-over-IP applications such as Skype, up to 90 percent idle time has been observed.

These measurements highlight the drawbacks of the current static spectrum allotment policy. There has been an exciting flurry of activities in engineering, economics, and regulation communities in searching for dynamic spectrum access strategies for improved spectrum efficiency. Various approaches have been proposed and studied. A taxonomy of dynamic spectrum access strategies can be found in [3].

In this article we focus on the overlay approach under the hierarchical access model of dynamic spectrum access [3]. Spectrum overlay was first envisioned by Mitola [4] and then investigated in the Defense Advanced Research Pro-

jects Agency (DARPA) Next Generation (XG) program as *opportunistic spectrum access* (OSA). The idea is to exploit instantaneous spectrum availability by opening licensed spectrum to secondary users. It directly targets spatial and temporal spectrum white space by allowing secondary users to identify and exploit local and instantaneous spectrum availability in a nonintrusive manner. Even in unlicensed bands, OSA may be of considerable value for spectrum efficiency (e.g., by adopting a hierarchical pricing structure to support both subscribers and opportunistic users).

To realize these potentials, many complex issues in technical, economical, as well as regulatory aspects need to be addressed. In this article we focus on technical aspects of OSA. We identify basic components, fundamental trade-offs, and practical constraints, and discuss open problems and recent advances. Based on the theory of partially observable Markov decision processes (POMDPs), we develop a decision-theoretic framework that leads to an optimal joint design of OSA, and a systematic examination of the interaction between signal processing for opportunity identification and networking for opportunity exploitation.

TECHNICAL CHALLENGES AND DESIGN TRADE-OFFS

While conceptually simple, OSA presents challenges not present in conventional wired or wireless networks. To protect spectrum licensees from interference while providing sufficient benefit to secondary users, OSA must rely on advanced signal processing techniques for instantaneous opportunity identification and sophisticated networking protocols for nonintrusive opportunity exploitation. The tension between the secondary users' desire for performance and the primary users' need for protection dictates the interaction between opportunity identification and opportunity exploitation, and the optimal design of OSA calls for a cross-layer approach that integrates signal processing with networking.

Basic design components of OSA include a spectrum sensor at the physical layer for oppor-


tunity identification, a sensing policy at the medium access control (MAC) layer for real-time decisions about which channels in the spectrum to sense, and an access policy, also at the MAC layer, to determine whether to access based on the sensing outcome. These three components should be jointly designed to maximize the throughput of secondary users while limiting the interference to primary users.

SPECTRUM SENSOR: FALSE ALARM VS. MISS DETECTION

The spectrum sensor of a secondary user identifies spectrum opportunities by detecting the presence of primary signals (i.e., by performing a binary hypothesis test). Sensing errors are inevitable: false alarms occur when idle channels are detected as busy, and miss detections occur when busy channels are detected as idle. In the event of a false alarm, a spectrum opportunity is overlooked by the sensor, and eventually wasted if the access policy trusts the sensing outcome. On the other hand, miss detections may lead to collisions with primary users. While both types of sensing errors are undesirable, reducing the occurrence of one generally comes at the price of increasing the occurrence of the other. Consider, for example, an energy detector. Choosing a larger energy detection threshold reduces the probability of a false alarm but increases the probability of miss detection. The trade-off between false alarm and miss detection is thus an important issue and should be addressed by considering the impact of sensing errors on the MAC layer performance in terms of throughput and collision probability. On a more fundamental level, which criterion should be adopted in the design of the spectrum sensor, Bayes or Neyman-Pearson (NP)? If the former, how do we choose the risks? If the latter, how should we set the constraint on the probability of false alarm?

SENSING POLICY: GAINING IMMEDIATE ACCESS VS. GAINING INFORMATION FOR FUTURE USE

Due to hardware limitations and the energy cost of spectrum monitoring, a secondary user may not be able to sense all the channels in the spectrum simultaneously. A sensing policy is thus necessary for intelligent channel selection to track the rapidly varying spectrum opportunities. The purpose of the sensing policy is twofold: catch a spectrum opportunity for immediate access, and obtain statistical information on spectrum occupancy for better opportunity tracking in the future. A balance has to be reached between these two often conflicting objectives, and the trade-off should adapt to the bursty traffic and energy constraint of the secondary user. When the user has no data to transmit, is it worthwhile to continue spending energy on spectrum monitoring? If so, how should the sensing policy change given that immediate spectrum access is no longer necessary? Clearly, such decisions should be made by taking into account the accuracy and energy consumption characteristics of the spectrum sensor.



■ **Figure 1.** A wireless LAN traffic measurement during an active FTP session.

ACCESS POLICY: AGGRESSIVE VS. CONSERVATIVE


Based on the imperfect sensing outcomes given by the spectrum sensor, the secondary user needs to decide whether to access. The objective of the access policy is to minimize the chance of overlooking an opportunity without violating the constraint of being nonintrusive. Whether the secondary user should adopt an aggressive or a conservative access policy depends on the operating characteristics (probability of false alarm vs. probability of miss detection) of the spectrum sensor, and joint design of them is necessary for optimality.

The above discussion provides a glimpse into the design complexity of OSA in a dynamic network environment with fading, random traffic, energy constraints, and competing distributed users. Is the optimal joint design tractable? Even if we arrive at an optimal solution, will it be too complicated to implement and too sensitive to environmental changes to be useful?

A DECISION-THEORETIC FRAMEWORK BASED ON POMDP

As an initial attempt to address the technical challenges outlined above, we introduce a decision-theoretic framework. Based on the theory of POMDP, this framework integrates the three basic components of OSA, leading to an optimal joint design of signal processing algorithms for opportunity identification and networking protocols for opportunity exploitation.

POMDP often suffers from the curse of dimensionality. The constraint on interference to primary users further complicates the problem. We have shown that, surprisingly, the structure of OSA admits a separation principle that decouples the design of the sensing policy from that of the spectrum sensor and access policy. This separation principle reveals the optimality of the myopic approach to design of the spectrum sensor and access policy, leading to closed-form optimal solutions. Furthermore, the design of the sensing policy is reduced to an unconstrained POMDP, where optimality can be achieved with deterministic policies. These results suggest a favorable trade-off between optimality and complexity of the OSA design.



■ **Figure 2.** Illustration of the set of all feasible sensor operating points (the operating point (ϵ, δ) can be achieved by randomizing between the optimal NP detectors designed under the constraints that the false alarm probability is no larger than $\epsilon^{(1)}$ and $\epsilon^{(2)}$, respectively).

NETWORK MODEL

Consider a spectrum consisting of N channels, each with bandwidth B_i ($i = 1, \dots, N$). These N channels are allocated to a network of primary users who communicate according to a synchronous slot structure. The traffic statistics of the primary network are such that the occupancy of these N channels follows a discrete-time Markov process with 2^N states, where the state is defined as the availability (idle or busy) of each channel.

We consider a group of secondary users seeking spectrum opportunities in these N channels. We focus on an ad hoc network where secondary users sense and access the spectrum independently. In each slot, a secondary user chooses a set of channels to sense and a set of channels to access based on the sensing outcome. Limited by its hardware constraints and energy supply, a secondary user can sense no more than L_1 ($L_1 \leq N$) and access no more than L_2 ($L_2 \leq L_1$) channels in each slot. To simplify notations and illustrate the basic idea, we consider $L_1 = L_2 = 1$. The decision-theoretic framework presented in this article, however, applies to the general case.

BASIC DESIGN COMPONENTS

As noted earlier, OSA has three basic design components: a spectrum sensor, a sensing policy, and an access policy.

Spectrum Sensor — By performing a binary hypotheses test, the spectrum sensor detects the presence of primary users in a chosen channel. Referred to as the receiver operating characteristic (ROC), the probabilities of false alarm and miss detection (ϵ, δ) specify the performance of the spectrum sensor. For a given ϵ , the largest achievable probability of detection $P_{D,\max}(\epsilon)$ (or equivalently, the smallest

achievable probability of miss detection $\delta_{\min}(\epsilon) = 1 - P_{D,\max}(\epsilon)$) can be attained by the optimal NP detector with the constraint that the false alarm probability is no larger than ϵ , or by an optimal Bayesian detector with a suitable set of risks [5, Sec. 2.2.1]. As illustrated in Fig. 2, the best ROC curve $P_{D,\max}$ forms the upper boundary of the feasible set of operating points. We also note that every feasible operating point (ϵ, δ) lies on a line that connects two boundary points and hence can be achieved by randomizing between two optimal NP detectors with properly chosen constraints on the probability of false alarm [5, Sec. 2.2.2]. Therefore, the design of the spectrum sensor is reduced to the choice of the desired sensor operating point. In other words, our objective is to find, sequentially in each slot, the optimal sensor operating point (ϵ^*, δ^*) in the feasible set to achieve the best trade-off between false alarm and miss detection. Note that the optimal operating point may vary from slot to slot.

Sensing and Access Policies — The sensing policy decides, sequentially in each slot, which channel to sense, and the access policy determines whether to transmit based on the sensing outcome. When the secondary user accesses an idle channel, a reward is accrued in this slot (e.g., we can define reward as the number of bits delivered). On the other hand, a collision with primary users occurs when accessing a busy channel.

The joint design of OSA is to choose the sensing and access policies together with the sensor operating policy that specifies the operating point (ϵ, δ) in each slot. The objective is to maximize the total expected reward accumulated over time under the constraint that the probability of colliding with primary users is capped below a prescribed level.

CONSTRAINED POMDP FORMULATION

Due to partial spectrum monitoring and sensing errors, the internal state of the underlying Markov process that models spectrum occupancy cannot be fully observed. Considering the constraint on the collision probability, we can formulate the joint design of OSA as a constrained POMDP as detailed below.

Reward and Objective Functions — A natural definition of reward is the number of bits delivered. For example, the reward for accessing an idle channel a can be defined as the bandwidth of channel a ,

$$R = B_a.$$

In a fading environment, the reward may also depend on the random fading gain of channel a .

We can define the objective function as the expected total number of bits transmitted in T slots:

$$J \triangleq \mathbb{E} \left[\sum_{t=1}^T R(t) \right]. \quad (1)$$

Note that the reward $R(t)$ obtained in slot t depends on the sensing action (which channel to choose), the access action (whether to transmit), the sensor operating point $(\epsilon(t), \delta(t))$, and the

state of the underlying Markov process (channel availability) in slot t .

This objective function is particularly appropriate when the underlying Markovian model only holds for a small number of slots due to rapid variations in spectrum occupancy statistics. When the spectrum usage of primary users is relatively static, we can use the transmission rate averaged over an infinite horizon or the total discounted bits as the objective:

$$J = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T R(t) \right], \quad \text{or} \quad J = \mathbb{E} \left[\sum_{t=1}^{\infty} \eta^t R(t) \right],$$

where $0 < \eta < 1$ is the discount factor. The latter is more appropriate for delay-sensitive messages where transmissions in the future are less rewarding.

Constraint and Joint Design — The design constraint is on the interference to primary users. Let ζ denote the maximum probability of collision allowed in any channel and any slot. Using the objective function defined in Eq. 1, we can formulate the joint design of OSA as finding the optimal sensor operating policy π_δ^* , the optimal sensing policy π_s^* , and the optimal access policy π_c^* given by

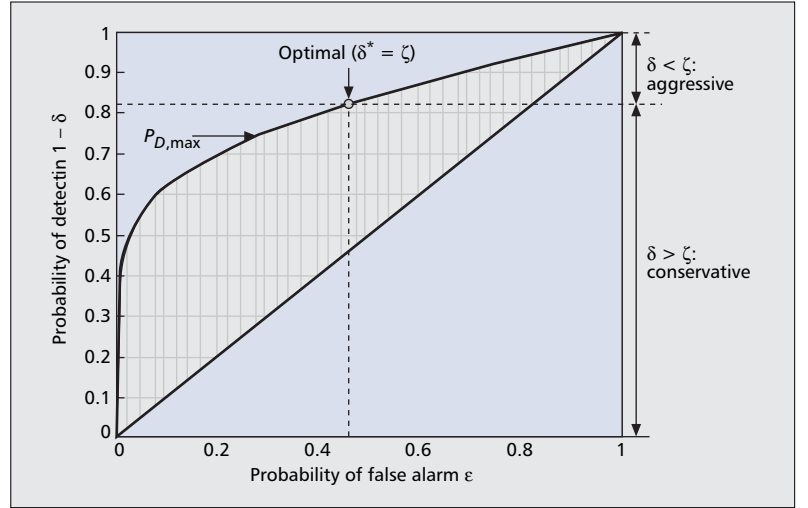
$$\{\pi_\delta^*, \pi_s^*, \pi_c^*\} = \arg \max_{\pi_\delta, \pi_s, \pi_c} \mathbb{E} \left[\sum_{t=1}^T R(t) \right], \quad (2)$$

subject to $P_c \leq \zeta$,

where P_c is the probability of collision determined by the chosen $\{\pi_\delta, \pi_s, \pi_c\}$.

Sufficient Statistic — The key to choosing the optimal actions in a given slot is the knowledge of the current state of the underlying Markov process. While the system state cannot be directly observed, the user can infer it from its decision and observation history. As shown in [6], the statistical information about the system state provided by the entire decision and observation history can be encapsulated in a belief vector $\Lambda(t) = [\lambda_1(t), \dots, \lambda_{2^N}(t)]$, where $\lambda_j(t)$ is the conditional probability (given the decision and observation history) that the system state is j at the beginning of slot t . Smallwood and Sondik have shown that this belief vector is a sufficient statistic [6]. Thus, a sensor operating policy π_δ defines the mapping from the current belief vector $\Lambda(t)$ to the sensor operating point $(\epsilon(t), \delta(t))$ used in this slot. Similarly, a sensing policy π_s maps $\Lambda(t)$ to the index of the channel to be sensed in this slot, and an access policy π_c maps $\Lambda(t)$ and the sensing outcome to the access decision. With a finite horizon T , the optimal policies are usually nonstationary; that is, the mapping from $\Lambda(t)$ to actions varies with time.

For a constrained POMDP (as we have here), we often need to resort to randomized policies to achieve optimality. In this case, π_s determines the probability of choosing each channel, π_c the transmission probability, and π_δ the probability density function of (ϵ, δ) . Due to the continuous action space, randomized policies are computationally prohibitive and implementationally cumbersome. Fortunately, as described below, the



■ **Figure 3.** An illustration of the interaction between the PHY layer spectrum sensor and the MAC layer access strategy (ϵ : probability of false alarm, δ : probability of miss detection, ζ : maximum allowable collision probability).

structure of the problem admits a separation principle that leads to deterministic policies without sacrificing optimality.

OPTIMAL JOINT DESIGN AND SEPARATION PRINCIPLE

We have established a separation principle for the joint design of OSA that provides a simple and explicit optimal solution to a seemingly intractable problem [7]. We have shown that the joint design can be carried out in two steps without losing optimality:

- Obtain the optimal sensor operating policy π_δ^* and the optimal access policy π_c^* by maximizing the instantaneous reward $R(t)$ in the current slot under the collision constraint.
- Obtain the optimal sensing policy π_s^* to maximize the objective function J given in Eq. 1 using π_δ^* and π_c^* obtained in the first step.

The separation principle decouples the design of the sensing policy from that of the spectrum sensor and access policies. As a consequence, the design of the sensing policy is reduced to an *unconstrained* POMDP, where optimality is achieved with *deterministic* policies. Furthermore, it reveals that the optimal sensor operating policy π_δ^* and the optimal access policy π_c^* can be obtained from a myopic approach that focuses solely on the instantaneous reward and ignores the impact of the current actions on the future reward. The joint design of π_δ^* and π_c^* is thus reduced to a static optimization problem with a simple, time-invariant, and closed-form solution. This closed-form optimal design of π_δ^* and π_c^* also allows us to quantitatively characterize the interaction between the physical (PHY) layer spectrum sensor and the MAC layer access strategy.

Optimal Spectrum Sensor and Access Policy in Closed-Form — As illustrated in Fig. 3, the set of feasible sensor operating points is partitioned into two regions by the maximum allowable collision probability ζ : the “conservative” region ($\delta > \zeta$) and the “aggressive” region ($\delta < \zeta$). When the

The alternative to a foresighted planning is the myopic approach that aims solely at maximizing the immediate reward. As revealed by the separation principle, the myopic approach to the design of the spectrum sensor and the access policy leads to the optimal solution.

sensor operates at $\delta > \zeta$, there is a high chance that a busy channel is detected as idle. This suggests that the access policy should be conservative to ensure that the collision probability is capped below ζ . Indeed, as shown in [7], when the channel is detected as busy, the user should always refrain from transmission; even when the channel is detected as available, it should only transmit with probability $\zeta/\delta < 1$.

On the other hand, in the region where $\delta < \zeta$, false alarms are more likely to happen. To reduce overlooked opportunities, the user should adopt an aggressive access policy: when the channel is detected as available, always transmit; even when the channel is detected as busy, still transmit with probability $(\zeta - \delta)/(1 - \delta) > 0$.

When the sensor operates at $\delta = \zeta$, the optimal access policy simply trusts the sensor: access if and only if the channel is detected as available. In other words, the access policy does not need to be conservative or aggressive to balance the occurrence of false alarms and miss detections. Note that at this point the access policy becomes deterministic. Interestingly, the optimal joint design of OSA defined in Eq. 2 requires that the sensor operate at this transition point $\delta^* = \zeta$ on the best ROC curve $P_{D,\max}$ in each slot, independent of the belief vector [7]. As a consequence, the optimal policies π_s^* , π_c^* are all deterministic.

The separation principle allows us to obtain, in closed form, the optimal access policy for any feasible spectrum sensor, as well as the optimal joint design. Extensions of the separation principle to the multichannel sensing case can be found in [8].

Low-Complexity Design of the Sensing Policy — We consider now the optimal sensing policy. This is a standard unconstrained POMDP to which solutions can be found in [6]. Our focus here is complexity reduction by exploiting the underlying structure of OSA.

An analysis given in [9] shows that the computational complexity of obtaining the optimal sensing policy is $\mathcal{O}(N^T)$, which grows *exponentially* with the horizon length T . The complexity mainly comes from the dimension 2^N of the sufficient statistic Λ , the foresighted planning for maximizing the overall throughput, and the continuously growing observation history. To achieve a favorable trade-off between optimality and complexity, we explore the possibility of circumventing each of these three sources of high complexity.

It has been shown in [10] that when channels evolve independently, we can find a sufficient statistic whose dimension grows linearly instead of exponentially with the number N of channels. Specifically, let $\Omega = [\omega_1, \dots, \omega_N]$, where ω_i is the (marginal) conditional probability that channel i is available at the beginning of a slot. Ω is a sufficient statistic if the channels are independent. This result points to the possibility of significantly reducing the computation and storage complexity of the optimal sensing policy.

The alternative to foresighted planning is the myopic approach that aims solely at maximizing the immediate reward. As revealed by the separation principle, a myopic approach to the design of the spectrum sensor and access policy leads to

the optimal solution. A myopic sensing policy, unfortunately, is generally suboptimal. An interesting finding is that when channels evolve as independent and identical Markov processes, a myopic approach is also optimal for the design of the sensing policy; we no longer need to trade immediate spectrum access for spectrum occupancy information [11]. Furthermore, a myopic sensing policy has a simple structure; selecting channels in each slot is reduced to a counting procedure. The secondary user only needs to set up pointers indicating the channels to which the last visits occurred most recently or the longest time ago [11].

The key to truncating the observation history without decimating performance is to exploit the *mixing time* of the underlying Markov process [9]. The mixing time quantifies how long it takes for the Markov process to approach its stationary distribution. When the Markovian dynamics of the spectrum occupancy have a mixing time of M , sensing outcomes obtained more than M slots ago provide little information on the current channel state. We can thus truncate the observation history to M slots, and the sufficient statistic Ω takes only a small number of values. Thus, the computational complexity of the optimal sensing policy is reduced from $\mathcal{O}(N^T)$ to $\mathcal{O}(N^M T)$, which is linear, rather than exponential, in the horizon length T . More important, this result suggests a systematic way of trading off performance with complexity by choosing an appropriate truncation parameter M .

OPEN PROBLEMS

The decision-theoretic framework presented here captures the fundamental design trade-offs in OSA: false alarms vs. miss detections of the spectrum sensor, aggressiveness vs. conservativeness of the access policy, and gaining spectrum access vs. gaining spectrum information in the sensing strategy. Many problems in both fundamental theories and practical implementations, however, remain open.

THEORETICAL ASPECTS

We have assumed that the transition probabilities of the underlying Markov process that models spectrum occupancy are known or have been learned accurately. Simulation results suggest that OSA design under this POMDP framework is robust to model mismatch [7]. OSA with an unknown Markov model is an interesting yet nontrivial problem. With an unknown model, a secondary user learns a good policy by comparing the observation and action trajectories under different policies and correlating rewards with actions. Formulations and algorithms for POMDP with an unknown model exist in the literature [12]. They provide useful tools for solving this problem.

The results on the low-complexity design of the sensing policy apply only to independent channels. Generalizations to systems consisting of dependent channels remain open. Furthermore, the robustness of the optimal design to mismatched and time-varying spectrum occupancy models needs in-depth investigation. Answers to these questions will establish the fundamental

trade-off across optimality, complexity, and robustness of this framework.

Energy constraints can further enrich the problem. The cost in each slot consists of the energy consumed in both sensing and transmission over fading channels. The design objective is to maximize the number of bits transmitted during the battery lifetime of a user subject to a constraint on the probability of collision. Under the energy constraint, the user may choose not to transmit when the available channel suffers from severe fading, leading to protocols that are opportunistic in both time and spectrum. The user may even skip sensing when the current belief vector indicates that no channel is likely to be available. Preliminary results on energy-constrained OSA in a fading environment can be found in [13].

Also of interest are cooperating schemes where secondary users sense and share partial spectrum maps [14]. Challenges here include characterization of the overhead associated with cooperation and the design of optimal policies.

PROTOCOL IMPLEMENTATION ASPECTS

We have not considered protocol implementation specifics in a general multihop ad hoc network with competing secondary users. In a general network the state of spectrum occupancy can be location-dependent; a channel available at a transmitter may not be available at the corresponding receiver. Furthermore, the ability to deal with hidden and exposed terminals and collisions among secondary users is crucial to the efficiency of OSA. Transceiver synchronization is also an important issue. In the presence of collisions and sensing errors, ensuring that a secondary user and its intended receiver hop synchronously in the spectrum with minimal control message exchange is a challenge not present in conventional MAC design.

An initial attempt at addressing the above issues can be found in [10]. Many questions, however, remain unanswered. How can we further reduce collisions among secondary users caused by hidden terminals and wasted spectrum opportunities caused by exposed terminals? Are classic collision avoidance schemes such as busy tone and dual busy tone feasible for OSA where we may not have a dedicated channel for the transmission of busy tones? What is the optimal power control for multihop ad hoc OSA networks? Since power control determines the area within which primary users may be affected by a particular secondary user, how do we choose the transmission power of secondary users based on that of primary users? How do the maximum allowable collision probability, channel fading, and sensing errors affect power control? Existing techniques for conventional ad hoc networks may inspire new ideas to address these unique challenges in OSA.

CONCLUDING REMARKS

In this article we have outlined some of the technical challenges of OSA and made an initial attempt at establishing a theoretical framework within which these challenges can be systematically and collectively addressed. We conclude

this article with a brief overview of strategic applications envisioned for OSA, and exciting research activities in the communications and networking communities. The former sketches some of the many promises of OSA, the latter our engineers' answers to whether these promises will be fulfilled.

POTENTIAL APPLICATIONS

Both commercial and military applications of OSA have been envisioned. Consider, for example, sensor networks deployed for carbon monoxide or traffic monitoring in metropolitan areas, opportunistic WiFi users at airports, or military units penetrating deep in unknown territory.

OSA presents an attractive approach to rapid deployment crucial to applications for disaster relief and emergency response. As an example, consider a disaster relief scenario where multiple rescue teams from different agencies and states may come together. The composition of such teams is likely to dynamically change through the course of the rescue effort. A related example is that of a multinational coalition force that may be involved in full-scale military operations, peacekeeping, and humanitarian relief operations in spatially contiguous areas. Such a force will probably rely on multiple sensor networks, some of which may be deployed as needed, to provide actionable intelligence. When the tempo of operations is high, it would be difficult and, even if possible, wasteful, to pre-allocate spectrum resources to the various actors and agents.

Tactical wireless networks are closed-loop systems with delay, partial models, and inaccurate knowledge of various parameters. As a consequence, they fall naturally under the purview of partially observable Markov decision processes we have discussed. Elements of the network must sense, decide, and actuate. For such a complex combat system with heavy traffic, large scale, and heterogeneous wireless devices, OSA may be the key to integrated sensing, communication, and actuation.

RELATED WORK

In this article we have mainly focused on the exploitation of temporal spectrum opportunities resulting from the bursty traffic of primary users. There is also a growing body of literature focusing on spatial spectrum opportunities that are static or slowly varying in time. Example applications include the reuse of certain TV bands that are not used for TV broadcast in a particular region. Due to the slow temporal variation of spectrum occupancy, opportunity identification is not as critical a component in this class of applications, and existing work along this line often assumes perfect knowledge of spectrum opportunities in the whole spectrum at any location.

At the physical layer, opportunity identification in the presence of fading and noise uncertainties has been studied [15]. Cognitive radio, the physical platform of OSA, has also received increasing attention recently. Spectrum monitoring testbeds [2] and cognitive radio prototypes [16] are being developed by researchers from both academia and industry. They validate the feasibility and practical value of theoretical

Both commercial and military applications of OSA have been envisioned. Consider, for example, sensor networks deployed for carbon monoxide or traffic monitoring in metropolitan areas, opportunistic WiFi users at airports, or military units penetrating deep in an unknown territory.

research and provide empirical data for spectrum occupancy modeling.

This list is by no means complete. For an overview of recent developments in OSA, readers are referred to [3], and to proceedings of workshops and conferences such as DySpan and CrownCom.

ACKNOWLEDGMENT

This work was supported by the U.S. Army Research Laboratory CTA on Communication and Networks under Grant DAAD19-01-2-0011, and by the National Science Foundation under Grants CNS-0627090 and ECS-0622200. The second author was supported in part by the DARPA Connectionless Networks Program.

REFERENCES

- [1] FCC Spectrum Policy Task Force, "Report of the Spectrum Efficiency Working Group," Nov. 2002, <http://www.fcc.gov/sptf/reports.html>
- [2] S. Geirhofer, L. Tong, and B. Sadler, "A Measurement-Based Model for Dynamic Spectrum Access in WLAN Channels," *Proc. IEEE MILCOM*, 2006.
- [3] Q. Zhao and B. Sadler, "A Survey of Dynamic Spectrum Access: Signal Processing, Networking, and Regulatory Policy," *IEEE Signal Processing*, May 2007.
- [4] J. Mitola, "Cognitive Radio for Flexible Mobile Multimedia Communications," *Proc. IEEE Int'l. Wksp. Mobile Multimedia Commun.*, 1999, pp. 3–10.
- [5] H. L. V. Trees, *Detection, Estimation, and Modulation Theory, Part I*, Wiley-Interscience, Sept., 2001.
- [6] R. Smallwood and E. Sondik, "The Optimal Control of Partially Observable Markov Processes over a Finite horizon," *Ops. Research*, 1971, pp. 1071–88.
- [7] Y. Chen, Q. Zhao, and A. Swami, "Joint Design and Separation Principle for Opportunistic Spectrum Access," *Proc. IEEE Asilomar Conf. Signals, Sys., and Comp.*, Oct. 2006.
- [8] Y. Chen, Q. Zhao, and A. Swami, "Joint Design and Separation Principle for Opportunistic Spectrum Access in the Presence of Sensing Errors," submitted to *IEEE Trans. In. Theory*, Feb. 2007; http://arxiv.org/PS_cache/cs/pdf/0702/0702158.pdf
- [9] D. Djonin, Q. Zhao, and V. Krishnamurthy, "Optimality and Complexity of Opportunistic Spectrum Access: A Truncated Markov Decision Process Formulation," *Proc. IEEE ICC*, June, 2007.

- [10] Q. Zhao et al., "Decentralized Cognitive MAC for Opportunistic Spectrum Access in Ad Hoc Networks: A POMDP Framework," *IEEE JSAC*, Special Issue on Adaptive, Spectrum Agile and Cognitive Wireless Networks, Apr. 2007.
- [11] Q. Zhao and B. Krishnamachari, "Structure and Optimality of Myopic Sensing for Opportunistic Spectrum Access," *Proc. IEEE Wksp. Toward Cognition in Wireless Networks*, June 2007.
- [12] D. Aberdeen, "A Survey of Approximate Methods for Solving Partially Observable Markov Decision Process," tech. rep., Nat'l. ICT Australia, Dec. 2003, <http://users.rsise.anu.edu.au/~daa/papers.html>
- [13] Y. Chen, Q. Zhao, and A. Swami, "Distributed cognitive MAC for energy-constrained opportunistic spectrum access," *Proc. IEEE MILCOM*, Oct. 2006.
- [14] C. Santivanez et al., "Opportunistic Spectrum Access: Challenges, Architecture, Protocols," *Proc. Wireless Internet Conf.*, Aug. 2006.
- [15] D. Cabric, S. M. Mishra, and R. W. Brodersen, "Implementation Issues in Spectrum Sensing for Cognitive Radios," *Proc. 38th. Asilomar Conf. Signals, Sys., and Comp.*, 2004, pp. 772–76.
- [16] J. Kleider et al., "Demonstrating Robust High Data Rate Capability on a Software Defined Radio Using Anti-Jam Wideband OFDM Waveforms," *Proc. MILCOM*, 2005.

BIOGRAPHIES

QING ZHAO [S'97, M'02](qzhao@ece.ucdavis.edu) received a Ph.D. degree in electrical engineering in 2001 from Cornell University. In 2004 she joined the Electrical and Communications Engineering Department at University of California at Davis, where she is currently an assistant professor. Her research interests are in the general area of signal processing, communications, and wireless networking. She received the 2000 IEEE Signal Processing Society Young Author Best Paper Award. She is an Associate Editor of *IEEE Transactions on Signal Processing* and an elected member of the Signal Processing for Communications Committee of the IEEE Signal Processing Society.

ANANTHRAM SWAMI is with the U.S. Army Research Laboratory where his work is in the broad areas of signal processing, wireless communications, and networking, including both sensor networks and MANETs. He was educated at the Indian Institute of Technology-Bombay, Rice University, and the University of Southern California, and has held positions with Unocal, USC, CS-3, and Malgudi Systems. He was a statistical consultant to the California Lottery, and developed a Matlab-based toolbox for non-Gaussian signal processing.