33rd Annual International Conference of the IEEE EMBS
Boston, Massachusetts USA, August 30 - September 3, 2011

CARER: Efficient Dynamic Sensing for Continuous Activity Monitoring

Lawrence K. Au, Alex A.T. Bui, Maxim A. Batalin, Xiaoyu Xu, William J. Kaiser

Abstract— **Advancement in wireless health sensor systems has triggered rapidly expanding research in continuous activity monitoring for chronic disease management or promotion and assessment of physical rehabilitation. Wireless motion sensing is increasingly important in treatments where remote collection of sensor measurements can provide an in-field objective evaluation of physical activity patterns. The well-known challenge of limited operating lifetime of energy-constrained wireless health sensor systems continues to present a primary limitation for these applications. This paper introduces CARER, a software system that supports a novel algorithm that exploits knowledge of context and dynamically schedules sensor measurement episodes within an energy consumption budget while ensuring classification accuracy. The sensor selection algorithm in the CARER system is based on Partially Observable Markov Decision Process (POMDP). The parameters for the POMDP algorithm can be obtained through standard maximum likelihood estimation. Sensor data are also collected from multiple locations of the subjects body, providing estimation of an individual's daily activity patterns.**

I. Introduction and Motivation

Continuous activity monitoring is becoming important for managing chronic disease and enabling physical rehabilitation [1]. Wireless health sensor systems play a critical role by providing objective measures of individuals' physical conditioning and functional capabilities. For example, recent research effort on activity monitoring with chronic obstructive pulmonary disease (COPD) focuses on methodologies that yield an accurate pattern of daily activity over time [2]. Similar research efforts are directed to monitoring individuals in neurological rehabilitation [3].

A long standing challenge for broad deployment of wireless health sensor systems is the extension of system operating lifetime under energy constraints. Performance of activity classification with sensor systems requires that wearable sensors provide continuous activation to enable knowledge of subject states. Advances appear in the energy-efficiency of microelectronic components including sensing, computation, and communication devices. However, these devices are now required in greater number on each sensor system with constantly advancing requirements in extending lifetime. The energy cost of sensor system operation is often dominant, in particular with the requirements for multi-modal sensor systems.

This work was supported by the National Institutes of Health/National Library of Medicine under Grant NLM T15 LM07356.

Lawrence K. Au, Maxim A. Batalin, Xiaoyu Xu, and William J. Kaiser are with the Electrical Engineering Department, University of California, Los Angeles, Los Angeles, CA 90095

Alex A.T. Bui is with the Department of Radiological Sciences, University of California, Los Angeles, Los Angeles, CA 90095

Context detection methods can be introduced to reduce energy consumption in sensor systems through the introduction of sensor activation scheduling compliant to context state [4], [5]. This approach has been based on empirical observations that activity states are largely repetitive and may display long periods when states are persistent. Specifically, by recognizing that successive activities tend to remain the same temporally, a measurement system may exploit *episodic* measurements to reduce unnecessary sensing and communication costs [6].

This paper introduces CARER (Continuous Activity Recognition with Embedded Reasoning), a software sensing support system that supports dynamic sensor scheduling algorithms by exploiting knowledge of contexts. CARER employs an architecture in which a mobile device acquires data from multiple wireless wearable sensors. The primary innovation of CARER is a new sensor operation scheduling algorithm, based on a Partially Observable Markov Decision Process (POMDP). This provides a direct method for reducing energy demand while ensuring classification accuracy. The algorithm stochastically determines if a new sensor measurement is necessary to estimate the activity in the next episode, and adjusts sensor usage accordingly. It is important to note that the CARER system makes real-time dynamic decisions as to whether to obtain additional sensor measurements for state classification. Such an approach allows real-time decision-making and intervention during the course of activity monitoring.

In this paper, the specific example of wearable sensors integrating Bluetooth technology are equipped with triaxial accelerometers. The mobile device supports the CARER algorithm and controls the acquisition of sensor data. In support of other applications, the CARER system includes user interfaces for development support. In particular, the graphical user interface of the CARER system is capable of displaying real-time sensor data, state classification output, and the overall system energy consumptions of the Bluetooth sensors.

II. System Architecture

CARER adopts the tiered architecture proposed in MEDIC [7]. The first tier consists of *wearable sensors* that can be placed on an individual for physical and physiological monitoring of biomedical signals. The second tier comprises a *mobile device* (e.g., smartphone) that is responsible for: (a) sensor data aggregation, (b) local data preprocessing, (c) intermediate data storage, (d) remote sensor control, and (e) relaying sensor data and local processing outputs to the remote server.

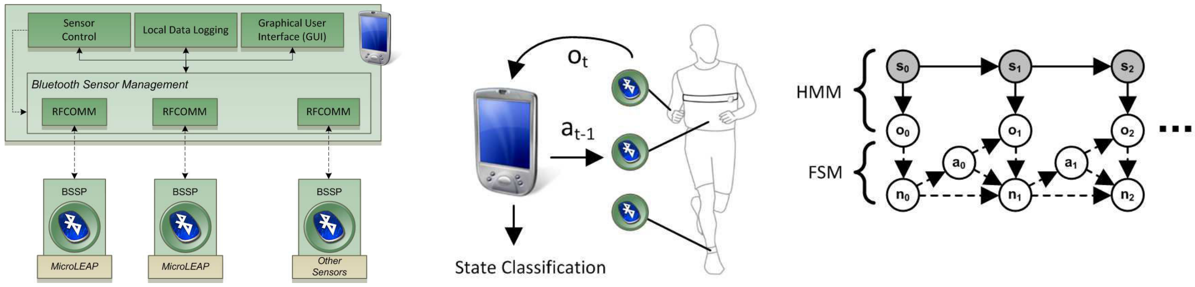

(a) System Architecture

(b) Sensing Methodology

Fig. 1: CARER System

A. Hardware System

As shown in Fig. 1a, the system consists of MicroLEAP wearable sensors ([8]) and a mobile device. The mobile device is responsible for establishing Bluetooth connections to gather sensor data and controlling the sensors. A system energy management unit on MicroLEAP is available for tracking total and average system energy consumption. An ASCII-based command protocol is implemented on MicroLEAP to facilitate serial data transfer and remote sensor control.

B. Software Architecture

1) Bluetooth Communication: Communication with Bluetooth sensors are performed via the RFCOMM interface in the underlying Bluez protocol stack [9]. Bluetooth sensors can be remotely controlled through the use of a pre-defined protocol. Each Bluetooth sensor occupies exactly one instance (thread) of the RFCOMM.

2) Bluetooth Sensor Management Unit: Multiple RFCOMM threads are instantiated by the Bluetooth Sensor Management Unit, which is responsible for creating and managing all Bluetooth connections between the mobile device and the sensors. Sensor data are identified with a unique sensor ID, and data are sent to upper-layer services and activities via messages. The unit is also responsible for monitoring the overall energy consumption of the Bluetooth sensors. MicroLEAP supports real-time system energy accounting; the traces of system energy consumption are sent periodically to the unit.

III. Methodology

The novel feature of the CARER system is its dynamic sensing algorithm based on a POMDP, which is a discrete-time, discrete-state control process that determines when to activate sensors for state classification [10]. The process consists of a hidden Markov model (HMM) that models the activity transitions, and a control policy based on finite-state machine (FSM) that controls the availability of Bluetooth sensors (see Fig. 1b). In a POMDP, the *State Space*, $|\mathbb{S}|$, contains the set of the activities of interest. The *Observation Space*, $|\Omega|$, contains the set of observation symbols extracted from sensor measurements. The *Action Space*, $|\mathbb{A}|$, consists of the set of all available sensor controls. (e.g., *activate* and *de-activate* Bluetooth sensors). The state-transition probability, $P(s_{t+1}|s_t)$, and observation probability, $P(o_t|s_t)$, have the same definitions as those in a standard HMM. For example, an HMM can model a process that includes *sitting* and *walking*. Accelerometers can be used as sensors, and features extracted from sensor measurements can be discretized.

The objective is to accurately detect activity patterns in different parts of the body given some resource constraint. At each discrete time step (or *sensing epoch*) t, sensor measurements are made if sensors are available (i.e., $a_{t-1} = Activate$). At the end of the epoch, an observation symbol o_t is generated, and it is used to compute the hidden state s_t. Simultaneously, the FSM maintains an internal state n_t and generates a new sensor decision a_t that determines whether sensors are to be activated next. POMDP estimates the posterior state probabilities based on the corresponding observation symbol.

A. Objective and Constraints

Both objective and constraint functions are specified by the *long-term average cost* for a given control policy π:

$$C_\pi = \lim_{T\to\infty} \frac{1}{T} \cdot \mathcal{E}_\pi \left[\sum_{t=0}^{T} c_\pi(s_t) \right] \tag{1}$$

Equation (1) takes into account the long-term effect of π, and it is independent of the initial state $s_{t=0}$ [10].

The objective function is defined as the misclassification cost of a one-step posterior estimation: $c(s_t, a_t) = 1 - \max_{s_{t+1}, o_{t+1}} P(s_{t+1}|o_{t+1}, s_t, a_t)$. The cost function is defined as the (*normalized*) energy consumption in one epoch: $e(s_t, a_t) \in [0, 1]$. The energy consumption can be measured by using the energy accounting capability on MicroLEAP or through offline measurements. The two functions are used to solve for the corresponding policy π in the POMDP problem; the objective (misclassification cost) and constraints (sensor energy) represent the performance-energy tradeoff.

B. Feature Extraction

In this paper, physical activity is defined as any bodily movement produced by skeletal muscles that results in energy expenditure beyond resting energy expenditure [11]. The Vector Magnitude Unit (VMU) is a commonly used feature for estimating energy expenditure in many accelerometry-based systems [2], [12]. Denote $f^{(i)}(t)$ as the VMU at time t; it corresponds to the magnitude of the applied force/acceleration measured from the i-th accelerometer: $f(t) = \sqrt{x^2(t) + y^2(t) + z^2(t)}$.

Denote the feature vector $\mathbf{f} = (f^{(1)}, ..., f^{(N)})$, where N is the number of accelerometers. The joint likelihood, $P(\mathbf{f}|s) = P(f^{(1)}, ..., f^{(N)}|s)$ is computed via supervised training: For each given state $s \in \mathbb{S}$, $f^{(i)}$ is discretized into two bins, and the boundaries for the discretization are determined in such a way that for each $s \in \mathbb{S}$ and some $\mathbf{f}'$, $P(\mathbf{f}'|s = j) = 1.0$ and $P(\mathbf{f}'|s \neq j) = 0.0$. In practice, the thresholds for the discretization are often determined experimentally by domain experts through specific training exercises [13]. Consequently, for each $S = j \in \mathbb{S}$ and $\mathbf{f}'$, $\mathbf{f}'$ is assigned a numerical observation symbol $O = i \in \Omega$:

$$P(O = i|S = j) = \begin{cases} 1.0 & \text{if } i = j \\ 0.0 & \text{otherwise} \end{cases} \quad (2)$$

In the feature extraction step, $P(o_t|s_t)$ takes on either zero or one for a given $s_t \in \mathbb{S}$. This implies full observability of the underlying state. Consequently, the state-transition probability, $P(s_{t+1}|s_t)$, can be obtained via maximum likelihood estimation using the training data. The sensing epoch is defined to be four seconds long at a sampling rate of 100 Hz (400 samples/epoch). One feature vector $\mathbf{f}$ is extracted in each epoch.

C. Classification

The posterior state probabilities are based on the Bayesian belief update equations [10]:

$$\begin{aligned} b_{t+1}(s_{t+1}) &= b_{t+1}(s_{t+1}|o_{t+1}, \mathbf{b}_t, a_t) \\ &= \frac{P(o_{t+1}|s_{t+1}, a_t) \cdot \sum_s P(s_{t+1}|s_t) \cdot b_t(s_t)}{P(o_{t+1}|\mathbf{b}_t, a_t)} \end{aligned} \quad (3)$$

where $\mathbf{b}_t = \big(b_t(0), ..., b_t(|\mathbb{S}| - 1)\big)$ is a column vector representing the state probabilities at time t, and $P(o_{t+1}|\mathbf{b}_t, a_t) = \sum_{s'} P(o_{t+1}|s', a_t) \cdot \sum_s P(s'|s) \cdot b_t(s)$ is the *conditional evidence* based upon the previous Bayesian update. Computing $\mathbf{b}_{t+1}$ is an iterative process that depends only on $\mathbf{b}_t$ (similar to the forward algorithm in HMM). The state classification output is defined as $\hat{s}_{t+1} = \arg\max_{s'} b_{t+1}(s')$.

D. Control Policy

The FSM provides the means for controlling sensors. An FSM is represented by $\pi = \{\mathbb{N}, \eta, \psi\}$, where $\mathbb{N}$ is the set of *nodes* $n \in \mathbb{N}$ in the FSM, $\eta = P(n_{t+1}|n_t, a_t, o_{t+1})$, and $\psi = P(a_t|n_t)$. η is responsible for maintaining the internal state information regarding the control policy, while ψ determines the relative frequency of different sensor controls $a \in \mathbb{A}$. The control policy is formulated as a constrained POMDP ([10], [14]), where the constraint is the average energy consumption per epoch. Equation (1) is used to specify these constraints.

Algorithm 1 Determining Sleep Time

```
1: INPUT: N_{t=0} = n_{t=0}, A_{t=0} = a_{t=0}, b_0, t = 0;
2: while (t < ∞) do
3:   if a_t == Activate then
4:     Acquire Sensor Data
5:     Extract Symbol : o_t
6:     Estimate and Classify: b_{t+1}
7:     n_{t+1} ~ P(N_{t+1}|n_t, a_t, o_{t+1})
8:     a_{t+1} ~ P(A_{t+1}|n_{t+1})
9:     ++t
10:  else
11:    T_sleep = 0
12:    repeat
13:      ++T_sleep
14:      Extract Symbol: o_{t+1} ~ U({0, 1, ..., |Ω| − 1})
15:      Estimate and Classify: b_{t+1}
16:      n_{t+1} ~ P(N_{t+1}|n_t, a_t, o_{t+1})
17:      a_{t+1} ~ P(A_{t+1}|n_{t+1})
18:      ++t
19:    until a_t == active
20:    Sleep for T_sleep epochs
21:  end if
22: end while
```

1) Dynamic Sensor Control Manager: One critical innovation in CARER is the development of the autonomous sensor control manager based on the dynamic sensing algorithm mentioned. The parameters of the policy π are implemented in a lookup table, with each value lies between zero and one (the probability space). As an example, Algorithm 1 shows an implementation of power-cycling Bluetooth sensors if $\mathbb{A} = \{Deactivate, Activate\}$. Lines 4-8 require one epoch of execution time. Lines 12-19 are used to determine the duration of sleep time. Because sensors are not available, the algorithm draws a random symbol (from the uniform distribution). Consequently, all probability distributions are known *a priori*, and each control evaluation step takes negligible (constant) amount of execution time relative to an epoch.

IV. Experimental Setup

Three MicroLEAPs are worn on the right wrist, waist, and right ankle to capture the activities of interest. The mobile device serves as the data aggregator, and continuously logs all sensor data for both training and testing. Three working days of training data (24+ hours) were collected from two subjects. Another working day of sensor data was collected for testing. The corresponding POMDP model parameters are as follows:

$\mathbb{S} = \{Static, RightArm, RightLeg, WholeBody\}$
$\mathbb{A} = \{Deactivate, Activate\}$
$\Omega = \{0, 1, ..., |\mathbb{S}| - 1\}$
$\pi = \{\mathbb{N}, \psi, \eta\}$
$P(s_{t+1}|s_t) = \frac{\text{count}(s_t, s_{t+1})}{\text{count}(s_t)}$
$P(o_{t+1}|s_{t+1}, a_t) = \begin{cases} P(o_{t+1}|s_{t+1}) & \text{if } a_t = Activate \\ 1/|\Omega| & \text{otherwise} \end{cases}$

V. Results

Fig. 2a and 2b show the predicted classifier accuracy as a function of the average energy consumption as obtained from the model. For instance, with both subjects, 40% of energy reduction can be achieved when maintaining an overall level

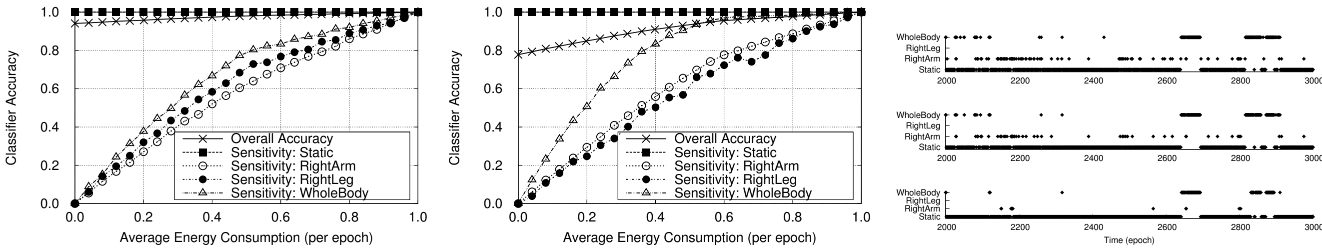

(a) Subject 1 (b) Subject 2 (c) Segment of Time-Series Data from Subject 1

Fig. 2: Classifier Accuracy vs. Average Energy Consumption

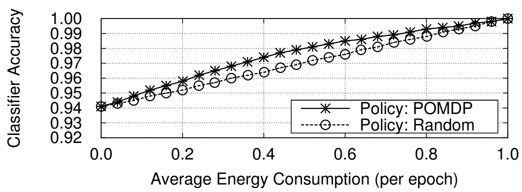

(a) Subject 1

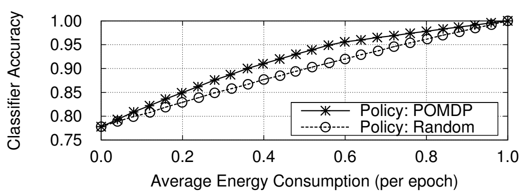

(b) Subject 2

Fig. 3: Policy Comparison

of 95% classification accuracy. The figures also show the sensitivity of each $s \in \mathbb{S}$ as a function of average energy consumption. *Sensitivity* is defined as the proportion of true positives that are classified correctly. Note that the proposed control policy performs more efficiently than traditional classification methods (e.g., decision-tree, naïve Bayes, etc.), because these methods assume sensors are available and active at all times. For instance, a naïve Bayes classifier would require continuous sensor activation (thus average energy consumption = 1.0 at all times) for state classification. The corresponding classification accuracy for naïve Bayes is thus linearly proportional to the average energy consumption, because the prior probabilities are assumed to be uniform. The POMDP approach improves the *overall classification accuracy* for a given energy consumption budget by using the state dependencies as additional information for classification. Fig. 2c illustrates a segment of the time-series as obtained by the control policy for different energy consumption constraints. The plots (from top to bottom) represent the ground truth, classification with average energy consumption = 0.6, and classification with average energy consumption = 0.2 respectively. The classification results approach the ground truth as the energy consumption budget increases; sensors are allocated correspondingly to detect different states.

Fig. 3 shows the relative performance of the FSM-based control policy against a completely random sensing policy (i.e., independent of observations and actions). The random sensing policy is generated according to the corresponding average energy consumption constraint. As expected, the FSM-based control policy consistently performs better than the random policy for both subjects; it allocates energy consumption for sensors more efficiently by considering $o \in \Omega$ and $a \in \mathbb{A}$ when determining the next control action.

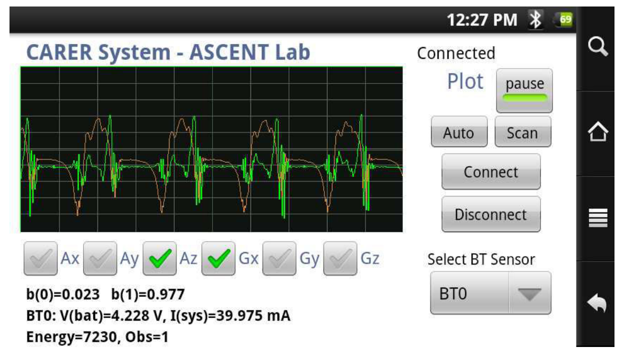

Fig. 4: CARER System: Screen Capture

Fig. 4 shows a sample screen capture from the CARER control system. It displays the sensor data, battery voltages and average system current consumption of the wearable sensors in real time.

VI. Conclusion

This paper demonstrates a new algorithm and system exploiting context state to enable continuous activity monitoring under an energy consumption budget. Specifically, by modeling the state dependencies, CARER dynamically selects the best episodes for sensor measurements. This dynamic management of sensor resources reduces the overall energy consumption of wireless health sensor systems.

The CARER system leads to significant energy reduction by recognizing the predictability of activities and exploiting the POMDP method with an approach that offers a solution compatible with processing technology available to mobile and wearable sensor systems. It is important to note that the CARER system provides a comprehensive profile of subject activity level; this offers an important future extension of CARER. Specifically, once a state is classified, other sensors and algorithms may then be recruited for the next stage of state analysis. In such a hierarchical configuration, sensors can be partitioned into different layers, with each layer providing different granularities. This will enable future applications of CARER that provide monitoring of subjects for determination of quality of motion for guidance of neurological rehabilitation (e.g., [3]) to accurate caloric energy expenditure for guidance of recovery interventions and promotion of health and wellness.

REFERENCES

[1] B. H. Dobkin, "Training and exercise to drive poststroke recovery," *Nat Clin Pract Neurol*, vol. 4, pp. 76–85, Feb 2008.

[2] A. Hecht, S. Ma, J. Porszasz, R. Casaburi, and for the COPD Clinical Research Network, "Methodology for using long-term accelerometry monitoring to describe daily activity patterns in copd," *COPD: Journal of Chronic Obstructive Pulmonary Disease*, vol. 6, no. 2, pp. 121–129, 2009. [Online]. Available: http://informahealthcare.com/doi/abs/10.1080/15412550902755044

[3] X. Xu, M. Batalin, and W. Kaiser, "Robust hierarchical system for classification of complex human mobility characteristics in the presence of neurological disorders," in *Body Sensor Networks 2011*, ser. BSN 2011, 2011.

[4] A. Krause, M. Ihmig, E. Rankin, D. Leong, S. Gupta, D. Siewiorek, A. Smailagic, M. Deisher, and U. Sengupta, "Trading off Prediction Accuracy and Power Consumption for Context-aware Wearable Computing," *Wearable Computers, 2005. Proceedings. Ninth IEEE International Symposium on*, pp. 20–26, Oct. 2005.

[5] Y. Wang, B. Krishnamachari, Q. Zhao, and M. Annavaram, "The tradeoff between energy efficiency and user state estimation accuracy in mobile sensing," in *In Proceedings of MobiCase*, 2009.

[6] L. Au, M. Batalin, T. Stathopoulos, A. Bui, and W. Kaiser, "Episodic sampling: Towards energy-efficient patient monitoring with wearable sensors," in *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE*, Sept 2009, pp. 6901 –6905.

[7] W. H. Wu, A. A. T. Bui, M. A. Batalin, L. K. Au, J. D. Binney, and W. J. Kaiser, "Medic: Medical embedded device for individualized care," *Artif. Intell. Med.*, vol. 42, no. 2, pp. 137–152, 2008.

[8] L. Au, W. Wu, M. Batalin, D. Mclntire, and W. Kaiser, "MicroLEAP: Energy-aware Wireless Sensor Platform for Biomedical Sensing Applications," in *Biomedical Circuits and Systems Conference, 2007. BIOCAS 2007. IEEE*, November 2007, pp. 158–162.

[9] Google Inc. Android Developers. [Online]. Available: http://developer.android.com/

[10] H. Yu and D. P. Bertsekas, "On near optimality of the set of finite-state controllers for average cost pomdp," *Math. Oper. Res.*, vol. 33, pp. 1–11, February 2008. [Online]. Available: http://portal.acm.org/citation.cfm?id=1527933.1527934

[11] P. D. Thompson, "Exercise and physical activity in the prevention and treatment of atherosclerotic cardiovascular disease," *Arterioscler. Thromb. Vasc. Biol.*, vol. 23, pp. 1319–1321, Aug 2003.

[12] B. G. Steele, B. Belza, K. Cain, C. Warms, J. Coppersmith, and J. Howard, "Bodies in motion: monitoring daily activity and exercise with motion sensors in people with chronic pulmonary disease," *J Rehabil Res Dev*, vol. 40, pp. 45–58, 2003.

[13] A. V. Rowlands, P. W. Thomas, R. G. Eston, and R. Topping, "Validation of the RT3 triaxial accelerometer for the assessment of physical activity," *Med Sci Sports Exerc*, vol. 36, pp. 518–524, Mar 2004.

[14] Q. Zhao, L. Tong, A. Swami, and Y. Chen, "Decentralized cognitive mac for opportunistic spectrum access in ad hoc networks: A pomdp framework," *Selected Areas in Communications, IEEE Journal on*, vol. 25, no. 3, pp. 589 –600, Apr. 2007.