Proc. of IEEE International Conference on Intelligent Robots and Systems, Nice, France, 2008.

Detecting and monitoring time-related abnormal events using a wireless sensor network and mobile robot

YuanYuan Li and Lynne E. Parker
Distributed Intelligence Laboratory, Department of Electrical Engineering and Computer Science
The University of Tennessee, 203 Claxton Complex, Knoxville, Tennessee 37996-3450
Email: {yli, parker}@eecs.utk.edu

***Abstract*— In this paper, we present an anomaly detection system that is able to detect time-related anomalies by using a wireless sensor network and a mobile robot. The sensor network uses an unsupervised fuzzy Adaptive Resonance Theory (ART) neural network to learn and detect intruders in a previously unknown environment. Upon the detection of an intruder, a mobile robot travels to the position where the intruder is detected to investigate by using its camera. The wireless sensor network uses a hierarchical communication/learning structure, where the mobile robot is the root node of the tree. Our fuzzy ART network is based on Kulakov and Davcev's implementation [8]. However, we enhance their work by extending the fuzzy ART neural network with a Markov model to learn a time series and detect time-related anomalies. Finally, a mobile robot is employed to verify whether the detected anomalies were caused by intruders. The proposed architecture is tested on physical hardware. Our results show that our enhanced detection system with mobile robot verification has a higher accuracy and lower false alarm rate than the original fuzzy ART system.**

I. Introduction

There are many advantages of using Wireless Sensor Networks (WSNs) to detect changes in the environment. Each individual node in the network can monitor its local region and communicate through a wireless channel with other nodes to collaboratively produce a high-level representation of the environment's states. By using such a network, large areas can be monitored with low cost.

In this research, we have investigated intruder detection in a previously By "unknown environment", we mean that the sensor signatures and types of anomalies are previously unknown to the WSN. We pre-deploy the static WSN into the environment. The sensor nodes first learn an initial model of the environment using a fuzzy ART neural network and a Markov model; we refer to this as the *normal* model of the environment. After the training period, any changes compared to the learned normal model are treated as anomalies possibly caused by an intruder. Upon the detection of the anomaly, an intrusion alert is generated, and an autonomous mobile robot responds to the alert by traveling to the place where the sensor nodes have detected the anomaly. A mobile robot response makes the system more flexible upon the detection of an intruder. The mobile robot is able to reach places and perform tasks that static sensors cannot. Mobile robots can also allow the design of a system where nodes can find power sources, request the dispatch of other nodes to perform tasks that require more sensing capability, and seek out repair.

The robot uses its additional sensors (e.g., a camera) to verify if there is an intruder in the area. We assume the intruder moves around in the area. Therefore, the mobile robot uses a camera to track moving objects. If there is a moving object, the robot declares that an intruder is detected.

An important challenge in WSN research is to determine a systematic procedure to train these networks so that they are sensitive only to real anomalies. To address this challenge, we have incorporated a machine learning technique into the WSN so that the networks can automatically learn to recognize normal and abnormal modes of operation. Our approach makes use of a fuzzy Adaptive Resonance Theory (ART) neural network, which was first implemented in [8]. The fuzzy ART neural network system is an unsupervised Artificial Neural Network (ANN) that can perform dimensionality reduction and pattern classification. The network can continually learn from new events without forgetting what has already been learned. No off-line training phase is required. The algorithm is simple enough to be implemented in the limited platform of the Crossbow motes [1], yet still achieve good performance.

However, a shortcoming of the original fuzzy ART approach is that it does not detect time-related changes. We have, therefore, enhanced the basic fuzzy ART system to enable it to learn a time series through the use of a Markov model. The approach builds a state transition model online during the initial period of deployment, and considers the built model as the normal model. After the training phase is over, any events that occur in the environment that do not fit the existing transition model are considered as abnormal events. Being able to model the expected sensor signatures for typical operations greatly simplifies the human designer's job; by enabling the system to autonomously characterize the expected sensor data streams, the network can learn the features of its environment that are important to monitor. This, in turn, allows the sensor network to perform autonomous anomaly detection to recognize when unexpected sensor signals are detected.

In a prior paper [11], we presented aspects of this proposed approach. This current paper, however, goes beyond this prior publication by presenting new results and analysis showing the benefit of our proposed approach over other competing techniques.

In this paper, we first review related work in Section II. Then, we present our approach in Section III. In Section IV, we present the hardware platform that we have used to test our system. Our experimental results from the physical experiments are presented in Section V. Finally, we summarize our findings in Section VI.

II. Related work

Our research objective is to design a scalable, efficient and robust anomaly detection system using WSN and mobile robots that will be deployed in an unknown environment. The desired characteristics of the learning algorithm are as follows:

1) *Able to detect anomalies in an unknown environment with minimum human supervision*. This characteristic makes supervised, offline learning algorithms, such as Bayesian networks [7], unsuitable for our application. Possible alternatives include Self-Organizing Maps (SOM) [3] and Adaptive Resonance Theory neural networks (ART) [8], as they are commonly-used unsupervised learning techniques.

2) *Able to easily scale to large numbers of motes*. Since a WSN typically has a large number of sensor motes, tuning the parameters of learning algorithms can be a long and tedious process. Therefore, the learning algorithm should have as few parameters to adjust as possible. The ART model allows the number of clusters to vary with problem size. Furthermore, it allows the operator to control the degree of similarity between members of the same cluster by means of the user-defined vigilance parameter.

3) *Able to support a hierarchical structure*. Heinzelman, et al., show in [6] that a hierarchical structure in a WSN is able to decrease communication requirements by reducing the size of the data transmitted; this in turn saves energy. The fuzzy ART learning technique works well with this hierarchical structure. Our sensor motes are able to run the same ART learning algorithm for cluster members as well as clusterheads, just with different inputs. Some machine learning techniques (e.g., association-rules [4]) might not be easily implemented as a single learning algorithm that works for different types of input.

4) *Computationally inexpensive*. The sensor motes generally have limited computational resources and limited power. Typically, machine learning techniques like Expectation-Maximization (EM) and gradient-based algorithms are computationally expensive.

5) *Memory efficient*. The learning algorithm has to be small enough to be implemented and installed on sensor motes with limited memory. Thus, learning algorithms that use particle filters might not be a good choice, since they require large amounts of memory during the learning process.

6) *Able to detect time-related anomalies online*. Much attention has been focused on time series analysis in WSNs. Many approaches detect time series anomalies at the communication level of the WSNs such as network traffic, package routing, radio channel selection, and so forth (e.g. [16]). Some works focus on predicting sensor values in order to improve the performance of data collection and reduce the communication effort, e.g. [15], [2]. While we have selected the fuzzy ART model for this research, the original approach could not detect time-related changes. Thus, in this paper, we present an algorithm to enhance the model so that it can detect time-related changes.

7) *Modular*. Our system is designed to be modular. Each component can be removed if its capability is not required. For example, if time series analysis is not of interest, then we can turn off the Markov model; the system would then simply detect anomalies in the sensor signatures.

8) *Able to continuously monitor the environment*. In our approach, human intervention is not required to reset the system after an anomaly has been detected. Our system can reset itself.

9) *Robust*. We use the robots' mobility to bring more coverage, sensing and processing capabilities to the WSN. There are some works that make use of mobile robots together with a WSN in other applications. For example, the authors in [9] explore parasitic mobility in WSNs. They propose a solution to the problems of power usage, node size, and node complexity in the form of parasitically actuated nodes. LaMarca, et al., [10] used robots to increase the feasibility of WSNs since sensor networks can acquire data but lack actuation, while robots have actuation but limited coverage in sensing. Schaffert [14] adapts sensor network models for use with information maps and verifies the ability of such maps to improve robot localization. Ren, et al., [13] focus on a fire and intruder detection application by using sensors only on a mobile robot.

10) *Able to adopt feedback*. A human operator or higher-level clusterhead may have a more accurate, higher-level view of the environment. We desire our system to be able to adopt its learning according to the feedback from this higher level. In the future, we plan to incorporate this feature in our system.

In summary, there is no evidence that a specific clustering algorithm performs better in all tasks or applications. However, some clustering techniques may be more suitable for some specific types of data or applications. We have chosen the fuzzy ART model [8] for our WSN implementation, because it satisfies most of our requirements for our applications. In future work, we will consider alternative approaches, such as SOMs, for the purposes of comparison. In this current paper, we present an enhancement to the original fuzzy ART model to detect time-related anomalies. To our best knowledge, no previous work addresses online intruder detection using a system that is able to detect time-related changes by using both a WSN and mobile robots.

III. Approach

In this section, we first introduce our network architecture. Then, we describe the basic fuzzy ART network. Subsection III-C then discusses our approach to incorporating a Markov model for time series analysis.

A. Architecture for the sensor networks

In our system, sensor motes are arranged hierarchically, as shown in Figure 1. In our WSN, sensor motes are divided into clusters. Each cluster has a clusterhead and multiple

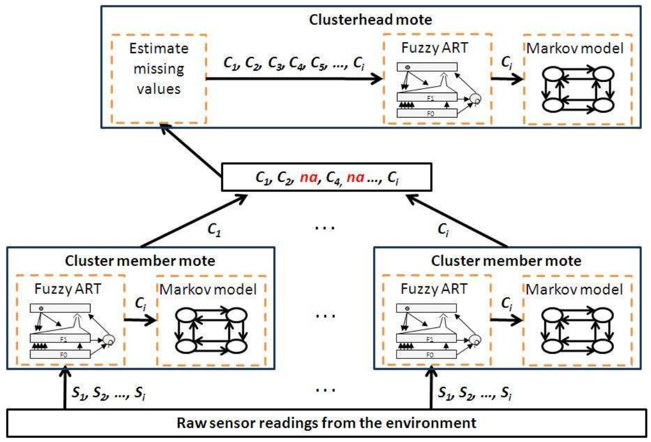

Fig. 1. Proposed fuzzy ART architecture, extended to estimate missing data and perform time series analysis.

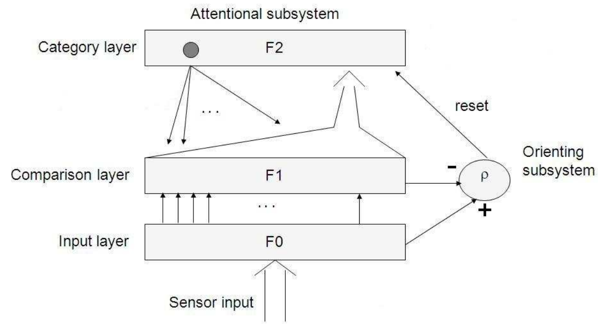

Fig. 2. A typical fuzzy ART architecture (see [8]).

cluster members. Each cluster covers a geometric region and is responsible for detecting the environmental changes in that region. Both cluster members and clusterheads run an identical detection system — a missing data estimator, a fuzzy ART network, and a Markov model. Cluster members read in raw sensor readings, s_i, (e.g., light and sound) from the environment as input, and then classify data into categories c_i. After the classification process, cluster members send their category labels to their clusterheads. The clusterheads first pre-process the collected category labels by identifying and estimating the missing values (through a process not described here; see our companion paper [12] for details). Then, the processed categorizations are used as input to their fuzzy ART neural network and are fused together to reduce the size of the data. The output of the fuzzy ART network is a category label c_i. After the classification process is finished, the system further checks if there are time-related changes. Clusterheads may have higher level clusterheads which classify their output class labels. Finally, the root mote obtains the final model of the environment. With this architecture, our WSN can be easily scaled to large numbers of sensors. At the same time, this hierarchical approach reduces communication, which in turn, saves energy in the WSN.

The fuzzy ART model alone cannot detect time-related abnormal events. For example, if people turning on the lights during the day and turning off the lights when they leave work is considered as a normal event, then an intruder only turning on the lights briefly in the evening should trigger an alarm. We enhance the basic fuzzy ART network by using a Markov model to detect these abnormal events. The Markov model takes the output category/state c_i from the fuzzy ART network and checks if the transition to the current state is probable based on the existing history. Note that the category c_i is the same as the Markov state. If the transition is not probable, an alarm is triggered. We believe that with this architecture, we can detect abnormal environmental changes as well as time-related changes. Our design is flexible because it allows the operator to turn off the Markov model if time series data is not of interest.

B. The fuzzy ART network

Kulakov and Davcev proposed a unsupervised fuzzy ART model for change detection in a WSN in [8]. Our basic fuzzy ART network is implemented in the same way. Figure 2 gives a representation of their fuzzy ART network. A typical fuzzy ART network has three layers: an input layer, a comparison layer and a category layer. The comparison layer takes an input vector from the input layer and transfers it to its best match in the category layer. If the best matching node is close enough to the input that is indicated by the vigilance parameter, the training process starts; otherwise, the next best matching node is selected, transformed and compared. If no existing nodes in the category layer meet the vigilance threshold, then a new node is generated and adjusted towards matching the input. The vigilance parameter influences the whole system: a higher vigilance level produces more categories, while a lower vigilance level produces fewer categories. For details on the fuzzy ART model, please refer to our companion paper [12].

C. Markov model extension

We enhance the existing fuzzy ART network by adding a Markov model to detect time-related changes. By definition, a Markov model is a discrete-time stochastic process with the Markov property, which states that, for a given process, knowledge of previous states is irrelevant in predicting the probability of subsequent states. At each time increment, the system may either stay in the same state, or transition to a new state. A Markov model is formally defined as a sequence of random variables $X_1, X_2, ...$, which, given the current state, the previous and next states are independent. Formally, $Pr(X_{n+1} = x | X_n = x_n, ..., X_1 = x_1) = Pr(X_{n+1} = x | X_n = x_n)$.

In a WSN setting, the Markov model is built during the training phase using the algorithm shown in Algorithm 1. Sensor motes periodically sense the environment at a fixed rate and feed the normalized sensor readings to the neural network to build categories of the environment. For each category/state (i), we keep an average time and the variance of the time the system remains in that particular state. Additionally, for each state we record the state transition probabilities, p_{ij}, to the next set of states. By doing so, an

Algorithm 1 Building the Markov model

for Each time step **do**
  if The current state is the same as the last time step **then**
    Record the time spent in this state.
  else
    Record the state transition.
  end if
end for
for Each state i **do**
  Find the mean μ_i and standard deviation σ_i of the time the system remains in state i.
  Find the transition probability p_{ij} for each possible state j.
end for

alarm will trigger if the amount of time in a state is either too short or too long. In a similar fashion, if a state transition is not probable, then this may also trigger an abnormal alarm. Thus, we can capture an anomaly from state transitions and from state occupancy time.

IV. Hardware platforms

Our wireless sensor network consists of static sensors (Crossbow motes) and mobile robots (Pioneer 3 robots). A Crossbow [1] mote contains a processing unit, a sensor module, and a communication module. The processing board contains an 8-bit processor at $8MHz$, a $128KB$ programming memory and a $512KB$ additional data flash memory. The wireless transmission range is around 10 meters inside a building. In our future work, we will extend the communication range by using intermediate motes as data routers. The sensor board has a buzzer, a light sensor, a microphone, 2 magnetometers and 2 accelerometers. For the experiments reported in this paper, we used the light and sound sensing components.

The mobile robot used in these experiments is a Pioneer 3 robot. Pioneer 3 is a mobile robot with a two-wheel differential drive. The mobile robot uses the Linux operating system and runs the Player-client/server device driver [5]. The robot uses a SICK LMS-200 range-finding laser for localization. The mobile robot can communicate with the sensor motes by having a mote attached to an MIB500 programming board through a serial connection. In our intrusion detection application, the robot runs the same fuzzy ART program as the motes. The robot takes the output from its cluster member motes and fuses them together to get the highest level representation of the environment. Thus, the mobile robot is a root mobile clusterhead with higher processing power and more sensing capabilities.

V. Experimental results

A. Intruder detection system

In order to detect abnormal events in a previously unknown environment, the sensor network first learns what is normal for the environment. Abnormal states of the environment are not kept in the sensor nodes due to memory limitations. Therefore, any events that do not match the existing normal model will be treated as abnormal events by the sensor motes. When an intruder is detected, a mobile robot moves to the area to investigate. We assume the robot knows the location of each cluster in advance. If the higher level clusterhead detects an anomaly (i.e., a category change after stabilization), the robot moves to the location of the cluster that detected the change. The mobile robot is the root clusterhead of the hierarchical fuzzy ART system.

In order to navigate in the environment, the mobile robot first creates a laser map using Simultaneous Localization and Mapping (SLAM). After an intruder has been detected by the sensor network, the mobile robot uses a wavefront path planning algorithm to plan a path from its current position to the goal position. During motion, it localizes itself using Monte Carlo localization.

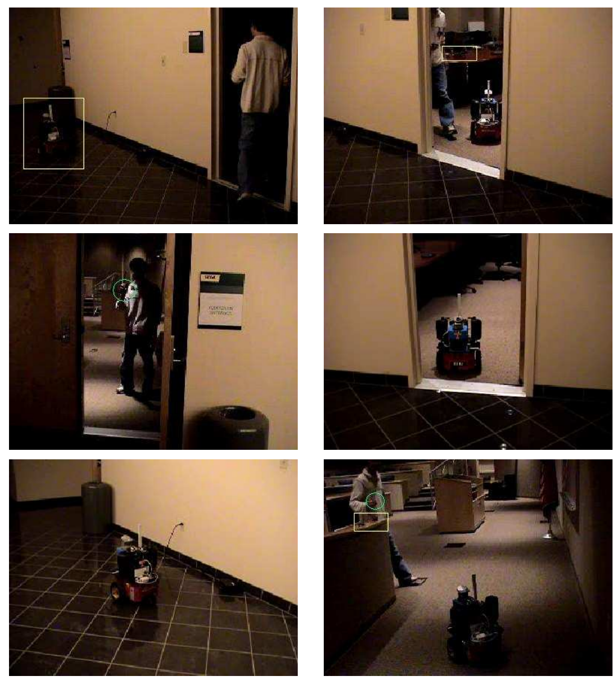

Fig. 3. Snapshots of the intruder detection system in operation at ORNL. Motes and the mobile robot are indicated by rectangles on the picture. The sound device carried by the intruder is indicated by a circle (read left to right, top to bottom).

We implemented and tested the intruder detection system on real motes along with a mobile robot and experimented with the system at both the University of Tennessee and Oak Ridge National Laboratory (ORNL). Figure 3 shows snapshots from the experiments at ORNL. We deployed 2 clusters of sensor motes in the environment. The first cluster was deployed into a conference room of ORNL's JICS building. The second cluster was deployed in an auditorium close by. The mobile robot was stationed in the hallway listening for abnormal changes. It detected abnormal changes by learning the combination of changes of the 2 clusterheads (sensor motes) deployed in the 2 rooms. The mobile robot ran the same learning algorithm as the sensor motes, namely, the fuzzy ART system. In the beginning, it was quiet and the lights were off in both rooms. The WSN learned that "quiet" and "dark" were normal in this environment. Then, an intruder entered the conference room and turned on the lights. The WSN detected the abnormal event and notified the robot. The robot planned a path using its wavefront path planner and moved to the conference room to check on the abnormal event — "light on". The intruder then moved to the auditorium. He turned on the lights and a buzzer to make noise in the auditorium. The robot detected the abnormal activities in the auditorium — "buzzer on" and "lights on". The robot then planned a path and moved to the auditorium to check on the abnormal event. Once the robot arrived at the auditorium, it used its camera to track the intruder.

In future work, we plan to remove the implementation dependency on synthetic intruder noises (i.e., buzzer).

B. Performance metrics

To evaluate our system, we collected statistics on the miss rate, false alarm rate, sensitivity and specificity. The miss rate is calculated as $\frac{FN}{(TP+FN)}$, where False Negative (FN) denotes the number of faults that the system failed to detect, and True Positive (TP) denotes the number of true faults that are detected by the system. The false alarm rate is defined as $\frac{FP}{(FP+TN)}$, where False Positive (FP) denotes the number of detected faults that were not true faults, and True Negative (TN) denotes number of "no faults" that were detected by the system. The sensitivity is defined as $\frac{TP}{(TP+FN)}$. The false alarm rate is defined as $\frac{FP}{(FP+TN)}$. Ideally, the values of sensitivity and specificity are at 100%, and the false alarm rate and miss rate are at 0%.

C. Temporal change detection experiment

In this experiment, we began by having the system learn the normal model; then, the testing began. Both training and testing were performed online. All sensors sampled the environment at a rate of 1 sample per second. Six motes were used during this experiment. One mote acted as a clusterhead, and the rest as cluster members of that mote. The cluster member motes were uniformly deployed around the clusterhead and all cluster members were within the communication range of the clusterhead. The vigilance levels for cluster members were set to 0.90, while those for the clusterheads were set to 0.97.

The training process took approximately 1.5 hours per trial. During the training period, states were visited multiple times. The averaged time was computed over the multiple visits of the same state. We treated this as a normal environment. Two sensors were used by cluster members — light and microphone. Raw light readings between 0 and 2000 indicated dark and light, respectively. Microphone readings came from a hardware detection system onboard. The values were binary — 1 indicates no noise is detected, and 0 indicates noise is detected. We used a buzzer as a sound source, which operates at $4Hz$. The sound sensor can detect the buzzer within a radius of 3 to 4 meters in our testing environment. Figure 4 is an example of the typical

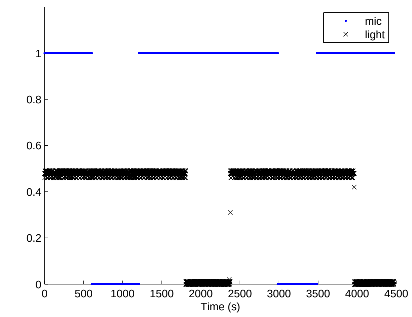

Fig. 4. An example: normalized light and microphone readings collected by a sensor mote. From time 0 to 510, light was on (0.5), microphone was on (1); From time 510 to 1200, light was on (0.5), microphone was on (0); From time 1200 to 1800, light was on (0.5), microphone was off (1); From time 1800 to 2300, light was off (0.1), microphone was off (1), etc.

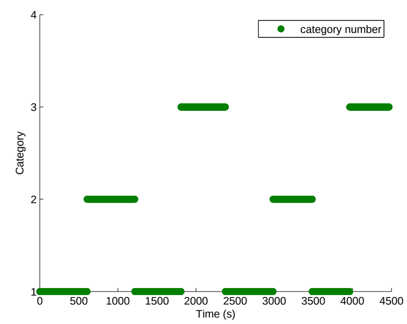

Fig. 5. An example of the detected changes from a cluster of 6 motes' fuzzy ART networks during the training phase from sensory data shown in Figure 4.

sensor readings collected from the environment. The light sensory readings are normalized between 0 and 1. Figure 5 shows the categories learned by the fuzzy ART neural network from the data shown in Figure 4 during the training period. After the classification process, a Markov model was built by using Algorithm 1. Figure 6 shows a Markov model built from the data shown in Figure 4 and Figure 5. Table I shows the mean and standard deviation values of the time the environment remained in each category/state before transiting to a different category/state. This is one of the trials of our experiment. In this particular example, the numbers on the Markov model is nicely rounded (e.g., 0.5 and 1). However, in more realistic situations, the Markov model can be much more complex; these experiments are designed to illustrate our approach.

TABLE I

TIME DURATION IN EACH STATE

Category	C1	C2	C3
Mean time (s)	571	555	538
Standard deviation (s)	62	75	43

Three different testing suites with four trials were run for each testing suite. In test suite 1, the environment started from "light and quiet" (category 1), and remained in that state

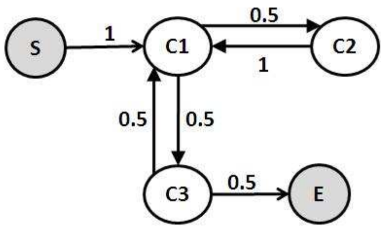

Fig. 6. An example of a learned Markov model for the training phase. The model was the normal model of the environment. States "S" and "E" denote the start and end states, respectively. They were manually added to the system. State C1 denotes lights were on and buzzer was off. State C2 denotes lights were on and buzzer was on. State C3 denotes lights were off and buzzer was off.

for 600 s. Then, it transited to "light and noisy" (category 2), and remained in that state for 600 s. Lastly, it transited to "dark and noisy" (category 4), and remained in that state for 600 s. Note that "dark and noisy" had never occurred before during our training phase. This was an abnormal event. This testing suite only contained a new abnormal state; however, it did not include any temporal-related changes.

In test suite 2, the environment started from "light and noisy" (category 2), and remained in that state for 600 s. Then, it transited to "dark and noisy" (category 4), and remained in that state for 300 s. Lastly, it transited to "dark and quiet" (category 3), and remained in that state for 600 s. The environment started with abnormal transitions to state 2, then the abnormal state 4 was detected. Lastly, an abnormal transition occurred from abnormal state 4 to state 3. This testing suite contained both abnormal events of a new abnormal state and abnormal time transitions.

In test suite 3, the environment started from "light and quiet" (category 1), and remained in that state for 300 s. Then, it transited to "dark and quiet" (category 3), and remained in that state for 900 s. The environment abnormally remained at state 1 too briefly and in state 3 for too long. This testing suite only contained time-related abnormal changes.

TABLE II

PERFORMANCE EVALUATION BETWEEN THE BASIC FUZZY ART AND ENHANCED FUZZY ART

		False Alarm	Miss Rate	Sensitivity	Specificity
Original fuzzy ART	mean	6%	59%	41%	94%
	stdev	12	4	38	12
Enhanced fuzzy ART	mean	6%	14%	86%	94%
	stdev	12	2	2	1

We used these testing suites to compare the performance of the basic fuzzy ART system (Kulakov and Davcev's implementation) and our enhanced fuzzy ART system. The experimental results are shown in Table II, which are averaged over 3 testing suites (for a total of 12 trials). Approximately 1500 observations were made from each sensor for each trial. The experimental results illustrate that our enhanced fuzzy ART system is able to detect more anomalies than the original fuzzy ART system (i.e., approximately 86% vs. 41%). This is due to the fact that our enhanced system learns a time series, and is able to detect time-related anomalies, whereas the original fuzzy ART cannot. Both the fuzzy ART system and the enhanced fuzzy ART system have a low false alarm rate (approximately 6%). To determine the significance of these results, we applied the Student's T-test to the miss rate and sensitivity results for the original fuzzy ART and our enhanced fuzzy ART. This test confirms that the differences in these results are statistically significant, with a confidence level of 99.5%. Thus, our enhanced fuzzy ART approach provides a significant improvement over the original fuzzy ART approach.

In these experiments, the time duration in each state is manually selected to illustrate the concept. In practical applications, the time duration could be very different, such as for comparisons between daytime versus nighttime expectations. However, in general, the proposed Markov model would be implemented similarly. We also realize that in some applications, the duration of time within a state is not of interest, but, instead, the order of the states is what is most important (e.g., if people always turn on the light before making noise in the room, regardless of the time duration in each state, then making noise in the dark room would be abnormal). In that case, the system would not have to maintain the time duration in each state. It can instead simply keep track of the expected state transitions.

D. *Intruder detection experiment*

After a change is detected in the environment, it does not necessarily mean that an intruder caused the anomaly. To determine if the anomaly is caused by an intruder, a mobile robot is sent to investigate using an additional sensor (i.e., a camera). In our enhanced intruder detection system, the sensor motes run our proposed change detection system, and a mobile robot serves as the mobile clusterhead. Once a change is detected by the mobile robot, it travels to the area and checks for an intruder using the camera mounted on the top of the robot. The camera tracks the intruder using a motion tracking program. The motion tracking program only detects moving objects. In the general case, we would want our mobile robot to carry its own light source or use a thermal image to detect a human in any lighting conditions, rather than just a moving object; this is the subject of future work.

If the mobile robot detects the intruder, the alarm is confirmed by the mobile robot. However, if the robot does not detect any intruder (human) within 120 seconds, it turns off the alarm and claims that there is no change in the environment. In this way, the robot does not miss any abnormal events occurring in the environment and at the same time reduces false alarms.

We ran the intruder detection system for the same sets of experiments in Section V-C except the abnormal state number 4 ("light off and noise") is not caused by an intruder. Instead, it is a normal state of the environment that never occurred in the initial learning process.

We compared the performance of the basic fuzzy ART system, the enhanced fuzzy ART system, and the enhanced

TABLE III

PERFORMANCE EVALUATION OF THE INTRUDER DETECTION SYSTEM

		False Alarm	Miss Rate	Sensitivity	Specificity
Original fuzzy ART	mean	46%	70%	30%	54%
	stdev	36	42	42	36
Enhanced fuzzy ART	mean	46%	17%	83%	54%
	stdev	36	20	25	40
Enhanced fuzzy ART w. mobile robot	mean	26%	17%	83%	74%
	stdev	19	25	24	19

fuzzy ART system with intelligent mobile robot responder. The results are shown in Table III, which are averaged over 3 testing suites for a total of 12 trials. We applied the Student's T-test to the miss rate and sensitivity results for the original fuzzy ART and our enhanced fuzzy ART. We also applied the Student's T-test to the false alarm rate and specificity for the enhanced fuzzy ART and the enhanced fuzzy ART with intelligent mobile robot responder. The tests confirmed that the differences in the results are statistically significant, with a confidence level of 99.5%. The experimental results illustrate that our enhanced fuzzy ART system and the enhanced system with mobile robot is able to detect more anomalies than the original fuzzy ART system (i.e., 83% vs. 30%). This is due to the fact that our enhanced system learns a time series and is able to detect time-related anomalies, whereas the original fuzzy ART cannot. The enhanced fuzzy ART system with intelligent mobile robot responder approach is able to reduce the false alarms compared to original fuzzy ART system and the enhanced fuzzy ART system (i.e., 26% vs. 46%). Thus, our enhanced fuzzy ART with mobile robot approach provides a significant improvement both in miss rate and false alarm rate over the original fuzzy ART approach.

We expect that if the mobile robot could provide feedback to the sensor motes regarding false alarms and the motes could correct their learning models based on this information, then the detection performance could be further improved. Additionally, the mobile robots could save their battery power by avoiding repeated checks of similar false alarms. Thus, in our future work, we plan to enhance our detection system by adding a feedback loop to the learning model, enabling learning from false alarms.

VI. CONCLUSION

We have presented an intruder detection system that is able to detect time-related anomalies by using a wireless sensor network and mobile robots. To our knowledge, this is the first intruder detection system that can detect time-related anomalies by using a sensor network to detect intruders and a mobile robot for traveling to the location where the intruder is detected. We have implemented and tested our system on physical motes and robots. The sensor network uses a fuzzy ART neural network to detect intruders. We have enhanced the original fuzzy ART system to detect time-related changes by using a Markov model. Our experimental results show that our detection system has high accuracy and is able to detect time-related changes. With an intelligent robot responding to alarms, the system is able to further reduce the false alarm rate.

As future work, we plan to enable the static sensors to make use of the mobile robot's feedback to improve the detection process. In addition, we are investigating how to use mobile robots to save energy in WSNs.

VII. ACKNOWLEDGMENTS

This research was sponsored in part by the ORNL SensorNet program. We also thank Michael Bailey for help with implementing the fuzzy ART algorithm, the operator control program, and integrating the Deluge system, and Shaddi Hasan for implementing the motion tracking program.

REFERENCES

[1] http://www.xbow.com/.

[2] Y. L. Borgne, S. Santini, and G. Bontempi. Adaptive model selection for time series prediction in wireless sensor networks. *Signal Processing*, 87(12):3010–3020, 2007.

[3] E. Catterall, K. Van Laerhoven, and M. Strohbach. Self organization in ad hoc sensor networks: An empirical study. In *The 8th International Conference on the Simulation and Synthesis of Living Systems*, 2002.

[4] Suan Khai Chong, Shonali Krishnaswamy, Seng Wai Loke, and Mohamed Medhat Gaben. Using association rules for energy conservation in wireless sensor networks. In *The 2008 ACM symposium on Applied computing (SAC)*, pages 971–975, Fortaleza, Ceara, Brazil, 2008.

[5] B. P. Gerkey, R. T. Vaughan, K. Stoy, A. Howard, G. S. Sukhatme, and M. J. Mataric. Most valuable player: a robot device server for distributed control. In *The IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 3, pages 1226–1231, 2001.

[6] W. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy-efficient communication protocols for wireless microsensor networks. In *International Conference on System Sciences*, Maui, HI, 2000.

[7] D. Janakiram, V.A. Reddy, and A.V.U.P. Kumar. Outlier detection in wireless sensor networks using bayesian belief networks. In *Communication System Software and Middleware (Comsware)*, pages 1–6, 2006.

[8] A. Kulakov and D. Davcev. Tracking of unusual events in wireless sensor networks based on artificial neural-network algorithms. In *Information Technology: Coding and Computing (ITCC)*, volume 2, pages 534–539, 2005.

[9] M. Laibowitz and J. A. Paradiso. Parasitic mobility for pervasive sensor networks. In *Pervasive Computing, Proceedings of the Third International Conference*, pages 255–278, Munich, Germany, 2005.

[10] A. LaMarca, W. Brunette, D. Koizumi, M. Lease, S. B. Sigurdsson, K. Sikorski, D. Fox, and G. Borriello. Plantcare: An investigation in practical ubiquitous systems. In *The 4th International Conference on Ubiquitous Computing*, pages 316–332, London, UK, 2002.

[11] Y. Li and L. E. Parker. Intruder detection using a wireless sensor network with an intelligent mobile robot response. In *IEEE Southeast Conference*, Huntsville, Alabama, April 2008.

[12] Y. Li and L. E. Parker. A spatial-temporal imputation technique for classification with missing data in a wireless sensor network. In *IEEE/RSJ 2008 International Conference on Intelligent Robots and Systems (IROS)*, Nice, France, September 2008.

[13] R. C. Luo, T. Y. Lin, H. C. Chen, and K. L. Su. Multisensor based security robot system for intelligent building. In *The IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, Heidelberg, Germany, September 2006.

[14] S. M. Schaffert. *Closing the Loop: Control and Robot Navigation in Wireless Sensor Networks*. PhD thesis, EECS Department, University of California, Berkeley, September 2006.

[15] D. Tulone and S. Madden. PAQ: Time series forecasting for approximate query answering in sensor networks. In *EWSN*, volume 3868, pages 21–37. Springer, 2006.

[16] Y. Zhang, Z. Ge, A. Greenberg, and M. Roughan. Network anomography. In *The 5th ACM SIGCOMM conference on Internet measurement*, pages 1–14, Berkeley, California, USA, 2005.