

Building Intrusion Detection with a Wireless Sensor Network*

Markus Wälchli and Torsten Braun

Institute of Computer Science and Applied Mathematics
University of Bern - Switzerland
{waelchli,braun}@iam.unibe.ch

Abstract. This paper addresses the detection and reporting of abnormal building access with a wireless sensor network. A common office room, offering space for two working persons, has been monitored with ten sensor nodes and a base station. The task of the system is to report suspicious office occupation such as office searching by thieves. On the other hand, normal office occupation should not throw alarms. In order to save energy for communication, the system provides all nodes with some adaptive short-term memory. Thus, a set of sensor activation patterns can be temporarily learned. The local memory is implemented as an Adaptive Resonance Theory (ART) neural network. Unknown event patterns detected on sensor node level are reported to the base station, where the system-wide anomaly detection is performed. The anomaly detector is lightweight and completely self-learning. The system can be run autonomously or it could be used as a triggering system to turn on an additional high-resolution system on demand. Our building monitoring system has proven to work reliably in different evaluated scenarios. Communication costs of up to 90% could be saved compared to a threshold-based approach without local memory.

1 Introduction

Subject of this paper is the detection and reporting of abnormal behavior in building monitoring. Conventional building monitoring systems address the problem by deploying video surveillance systems (see Section 2.1). Such systems have a number of drawbacks, though. The system is expensive in terms of hardware, storage, and communications. In particular if wireless technology is used, energy for communication becomes a critical issue. Furthermore, collecting multiple video streams imposes high demands on storage, online monitoring and video analysis. For example, the more video screens a security guard has to monitor, the higher is the probability that he misses some relevant information. Finally, a permanently active camera system is unpleasant for the office staff.

* The work presented in this paper was partly supported by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322.

The envisioned system involves numerous research questions: Is it possible to develop an energy-efficient reliable building surveillance system? Can the data volume be reduced while minimizing the probability of losing relevant information? Can privacy of staff be improved? Is the system economic?

Wireless sensor network technology provides means to develop such a system. Sensor networks are economic due to their lack of wires. Energy-efficiency is achieved by parallelism, distributed computing, and in-network processing. Decisions can be made in the network, decreasing communication load and storage requirements. Hence, office occupation can be processed and filtered in the network. Because no pictures are taken, the identity of office staff can be hidden. For these reasons, we provide a wireless sensor network for building monitoring. The system can run standalone or it can be used to trigger a high-resolution system such as a (wireless) video surveillance system. A hybrid system would optimize accuracy, while keeping costs low. Unless something critical is detected by the sensor network, video surveillance is turned off. During this time, energy and storage costs are minimal and the identity of the office staff is concealed.

Tininess, resource constraints, need for long-term operation, and dependency on batteries impose severe restrictions on wireless sensor networks. Hence, services provided in sensor networks need to be lightweight in terms of memory and processing power. Communication costs should be low. We satisfy these requirements by filtering normal (i.e., known) building occupation within the network. Normal behavior is temporarily learned by short term memory using an aging mechanism. The memory has been implemented as Adaptive Resonance Theory (ART) neural network. ART neural networks compress observed event patterns to a single value representing the ART decision (known | unknown). Unknown patterns are reported to the base station. At the base station, the reported local decisions are fed to a binary ART neural network that performs the system-wide anomaly detection. ART neural networks are adaptive and learn sequentially. They require low storage and communication costs. These features are advantageous for use together with sensor nodes.

Section 2 discusses related work in monitoring systems, event classification and anomaly detection. Section 3 presents the office monitoring system, including ART neural networks. Experiment setup and evaluation are described in Section 4. The paper ends with conclusions and an outlook in Section 5.

2 Related Work

Previous work has covered a wide variety of related topics including visual sensor networks, classification and anomaly detection in wireless sensor networks.

2.1 Visual Sensor Networks

Due to their potential, wireless video surveillance systems have recently gained a lot of attention in building or structural health monitoring. Much effort has been put into tailoring video streaming techniques to the requirements of resource-constraint systems. Efficient video coding [1] and reliable routing [2] have been

addressed. A dual-camera sensor network was proposed in [3]. A low-resolution camera is permanently used, whereas a high-resolution camera is only used on demand. Security and privacy issues have been investigated in [4].

In [5] an event-based triggering system was used for structural health monitoring of bridges. A wireless sensor network is used to control a video camera. If the wireless sensor network throws an abnormal event, the video camera is activated and zooms into the area of interest. The application goal of the system resembles our own. However, events are modeled based on thresholds. These thresholds have to be determined by system experts and work only for the specific deployment. In contrast, our system determines abnormal behavior in an unsupervised way. Moreover, the used memory provides pattern recognition.

2.2 Event Classification

Event classification approaches can mainly be divided into two categories. There are approaches that classify time-discrete events, while other approaches address continuous event patterns that evolve over time.

Advanced approaches that classify time-discrete events implement classification techniques from statistical and pattern recognition research fields. A wide range of statistical classifiers such as maximum likelihood (ML) estimators and support vector machines have been proposed in the SensIt project [6]. Apart from statistical methods, classifiers based on pattern recognition have been proposed. In [7] classification rules are learned by simulated annealing. Clustering and consensus have been used in [8] to learn and classify events in a fence monitoring application. In [9] a self-learning Fuzzy Logic Controller (FLC) learns and classifies different light sources. Fuzzy ART neural networks have been used in [10] to compress and classify multiple sensor readings on sensor node level.

Classifying events that evolve over time is complex. Hence, most proposed solutions for sensor networks focus on threshold-based systems [11], [12], [13]. [11] proposes a declarative query language to model events and event areas. [12] proposes a multi-tiered classifier based on decision trees. [13] proposes a hierarchical classifier, where threshold-based decisions are done on different layers. All these approaches require a priori knowledge of the expected events and thresholds. Context-aware and Time Delayed Neural Networks [14] have been used to classify bird songs. This approach requires periodic learning phases and specific neural networks for every kind of bird song. Finally, transferable belief states and particle filters have been used in [15]. The approach is similar to Hidden Markov Models (HMM), only the underlying physical process does not need to be random. Resource requirements are again high.

All presented approaches require pre-determination of expected events. This prevents any dynamic intrusion detection. However, in intrusion detection it is neither possible to determine all normal behavior nor to define any possible kind of intrusion in advance. Hence, online learning is required.

2.3 Anomaly Detection

Apart from event classification, anomaly detection systems have been proposed. As long as only distinction of abnormal from normal behavior is required, binary classifications are sufficient. In [16] spatiotemporal anomalies in gas distributions in underground coal mines trigger alarms. Bayesian networks are used. Missing alarms endanger life of miners. Hence, specific wired and power-supplied sensor networks are used. Anomaly and intrusion detection has gained much attention in computer security. Algorithms based on State Vector Machines (SVMs), Fuzzy Logic Controllers (FLC), Principal Component Analysis (PCA), Hidden Markov Models (HMMs), or Instance-based Learning (IBL), have been introduced [17]. Intrusion detection systems (IDS) have been proposed for security in wireless sensor networks [18], [19]. Rule-based voting is applied to prevent network attacks. [20] proposes an intruder detection system that combines Fuzzy ART mechanisms with Markov chains. Fuzzy ART neural networks are deployed on sensor nodes to guide a mobile robot. The approach requires a learning phase and loses learning capability thereafter. This makes any subsequent online adaptation and learning impossible. In contrast, our approach provides persistent learning based on short-term memory and is thus very communication and storage efficient.

Finally, anomaly detection in ad-hoc and sensor networks has been addressed by Artificial Immune Systems (AIS) [21], [22]. These systems show similar performance to ART neural networks, but are less compact and require more memory.

3 Office Monitoring

In this section the office monitoring application is introduced. We have used sensor nodes for a number of reasons. The system is lightweight, cost-effective, and easy to deploy. In particular no wires are required. The system conceals the identity of persons working in the office. Only sensor activation patterns can be determined. Thus, the system ensures privacy to the office staff. A high resolution system providing more detailed information could optionally be triggered. Thus, the application of the secondary, high resolution system could be restricted to periods of abnormal office occupation.

The deployment of the office monitoring system is depicted in Figure 1. In the current deployment the ART based anomaly detection software and the event detection and tracking functionality of DELTA were implemented. A common office room of approximately $26.5 m^2$ providing two working places was monitored. The room contains 5 office cabinets, 4 tables and two file cabinets. In total, 11 wireless sensor nodes were deployed: 3 nodes measuring Passive Infrared (PIR), 7 nodes measuring vibration and one node acting as gateway (base station). The 3 sensor nodes that measure PIR were placed such that the two working desks and the office entrance were monitored. The 7 vibration sensors monitored activation in the office cabinets and in the file cabinets. The sensor nodes do not report every sensor reading to the base station. This would be too communication intensive. Instead, series of sensor readings, 10 in the current implementation, are collected and processed on node level. The resulting local pattern classifications

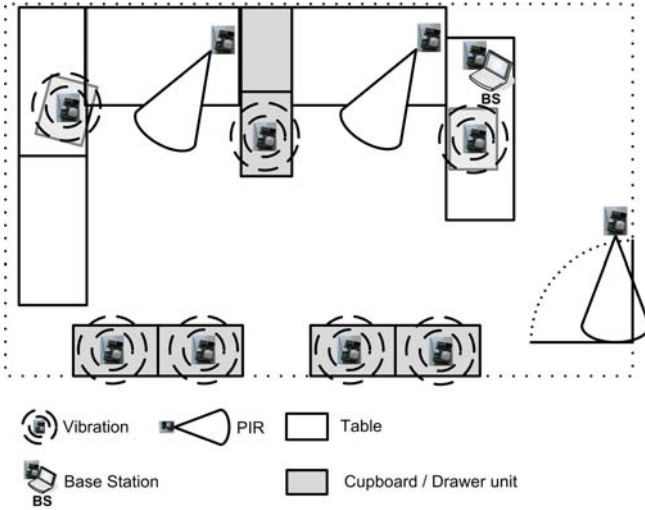


Fig. 1. Office monitoring deployment

(known | unknown) are then sent to the base station, where the system-wide anomaly detection is performed.

3.1 System Design

Our building monitoring system can be summarized as follows. Series of sensor readings are periodically collected and processed on node level. In the current implementation the PIR and vibration sensors are processed every 2 s. Since series of ten sensor readings are processed, i.e., the monitoring resolution on node level is 20 s, time vectors $\mathbf{x} = \{x_1, \dots, x_i; i = 10\}$ are locally processed by ART neural networks. As mentioned before, our ART neural networks implement an adaptive short term memory, which is based on an aging mechanism. The sensor nodes can store a certain number of prototype time vectors \mathbf{y} . Prototypes \mathbf{y} that are not matched by an input \mathbf{x} become older. When the memory of the ART neural network is full, the oldest prototype \mathbf{y} is replaced by the current input vector \mathbf{x} . With this mechanism learning can be continuously maintained. If \mathbf{x} is not recognized in the ART memory, i.e., the input pattern is unknown, a 1 is signaled to the base station. Otherwise, no report is sent. It is not necessary to report locally known patterns, i.e., to signal a 0, because the base station expects the presence of a known time vector \mathbf{x} if no report has been sent by the according node. Because each sensor node reports for every monitoring interval whether it has locally detected an unknown time vector \mathbf{x} or not, the base station is able to perform global decisions based on the collected reports. Since ten sensor nodes are deployed, the base station processes space vectors of the form $\mathbf{z} = \{z_1, \dots, z_j; j = 10\}$, where $z_i \in \{0, 1\}$ is the output of sensor node i .

An example may clarify the functionality. We assume, a person is searching two cabinet drawers that are equipped with vibration sensors 5 and 6. The first drawer is opened, then the second one, before the first drawer is closed, again followed by the second one. This is all done within a monitoring period of 20 s. Hence, time vector \mathbf{x}_5 containing the sensor readings of sensor node 5 might look like [0, 24, 12, 0, 0, 0, 14, 22, 0, 0], while \mathbf{x}_6 is [0, 0, 0, 33, 0, 0, 0, 0, 41, 13]. Values of 0 mean no activation of the according sensor (PIR or vibration). Both vectors are locally processed by Fuzzy ART neural networks. We assume that \mathbf{x}_5 is recognized by the memory of sensor node 5, while \mathbf{x}_6 is not recognized by the memory of sensor node 6. Since Fuzzy ART neural networks store prototypes, this recognition behavior is possible. The other sensor nodes have not measured activation in the current monitoring period, i.e., they recognize time vector [0, 0, 0, 0, 0, 0, 0, 0, 0, 0] in their memory. In this example, only sensor node 6 reports a 1 (unknown time vector). Accordingly, the binary space vector \mathbf{z} processed at the base station is [0, 0, 0, 0, 0, 1, 0, 0, 0, 0]. The base station implements a binary ART neural network to process \mathbf{z} . The output of the binary ART at base station is again 1 (unknown), if \mathbf{z} is not recognized and 0 otherwise.

ART neural networks are predestinated to meet storage, local pattern recognition, and filtering requirements. The storage costs of an ART neural network are determined by the size N of the input vector (i.e., of \mathbf{x} or \mathbf{z}) and the memory capacity M . M determines how many categories (represented as prototypes) of input vectors can be stored. Considering M categories and an input size of N , the algorithmic complexity of an ART neural network is $O(MN)$, both in terms of time and memory [23]. In our implementation we have used the ESB sensor node platform [24]. The available memory on these nodes for local signal processing is in the order of 300 bytes. Considering the available memory, the input vector size of 10, and 2 bytes to store floating points, 10 prototypes are stored in the memory of our ART neural networks. Finally, traditional ART neural networks return the category number if a category is determined for a given input \mathbf{I} and -1 otherwise. In contrast, our ART-based event detector returns 0 (known) if \mathbf{I} is recognized and 1 (unknown) otherwise. Hence, our ART neural networks provide the following features:

- ART neural networks are very lightweight.
- Known input vectors are filtered by the ART neural network, which saves communication costs.
- A data compression of N to 1 is achieved, which reduces data volume.

Next, the theory of ART neural networks and some adaptations to account for our system requirements are introduced.

3.2 ART Neural Networks

Adaptive Resonance Theory (ART) neural networks [25] represent a special kind of adaptive memory with sequential learning ability. Any present input

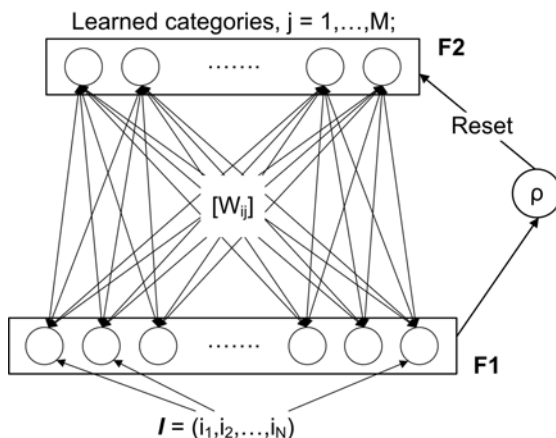


Fig. 2. ART neural network architecture

$\mathbf{x} = \{x_1, \dots, x_N\}$ is fed into the ART neural network. The present input is classified in respect to a number of stored prototypes, which represent learned classes of input patterns \mathbf{x} . If the present input can be classified, the respective prototype is updated. Otherwise, a new prototype is created unless the whole memory capacity is utilized. ART neural networks have been designed to process binary input patterns (binary ART) and analog input patterns (Fuzzy ART).

The architecture of any ART neural network is shown in Fig. 2. Any ART neural network is an unsupervised learning system. It consists of two layers, a comparison layer F1 with N neurons representing the attributes of a given input and a recognition layer F2 composed of M neurons representing the prototypes (categories). The weight matrix $W_{i,j}$ is the memory of the ART. The sensitivity ρ controls the recognition behavior of the ART neural network.

ART-based Event Recognition

input: Input vector I ;

output: Number representing category j to which I belongs;

begin

Compute similarity s_j to each prototype j in F2;

Sort the s_j in descending order;

for each s_j **do**

if $s_j > \rho$

Update the weights $W_{i,j} = I \cdot \alpha + W_{i,j} \cdot (1 - \alpha)$;

return 0;

end

if maximum number of categories is not reached

Commit uncommitted neuron n in F2;

return 1;

else

replace oldest category in F2;

return 1;

end

The operation of our ART-based event recognizer is described in the pseudocode above. The weight matrix $W_{i,j}$ is the memory of the ART. First, the similarity s_j between every prototype j and the input vector \mathbf{I} is determined. The Euclidean distance is used to determine similarity s_j between \mathbf{I} and any of the stored prototypes j . The resulting list of similarities is sorted in descending order. The ordering is necessary, because otherwise a prototype might be chosen (similarity applies), though a more similar prototype exists. This case could happen if more than one s_j exceeds the sensitivity threshold ρ . The resulting list of similarities is evaluated in respect to ρ . If an appropriate category is found, the weights of the according prototype are updated and 0 (known) is returned as classification output. If no category could be determined, a 1 is reported (unknown pattern).

The parameters of the ART neural network are explained in the following. If s_j exceeds sensitivity threshold ρ , \mathbf{I} is assigned to category j . A high value for ρ implies fine-grained memory (many small categories), since the input needs to match a category exactly. On the other hand, low values mean coarse recognition (few large categories). A second parameter that has impact on the behavior of the ART neural network is the learning rate α . If an input \mathbf{I} is assigned to a category j , the stored prototype j is updated according to the weighted sum of \mathbf{I} and j , i.e., $W_{i,j} = \mathbf{I} \cdot \alpha + W_{i,j} \cdot (1 - \alpha)$. Hence, the learning rate defines the weights α for the input and $(1 - \alpha)$ for the stored prototype. If α is high (e.g., 0.8), \mathbf{I} is weighted 0.8 and the stored prototype j is weighted 0.2. Accordingly, high learning rates reinforce the impact of the current input. Traditional ART neural networks return the category number if a category is determined for a given input \mathbf{I} and -1 otherwise. On the other hand, our ART-based event recognizers return 0 (known) if \mathbf{I} is recognized and 1 (unknown) otherwise.

Traditional ART neural networks provide a long-term memory of M categories. When all M categories in F2 are used, the learning capability of the ART neural network is exhausted. Any new input pattern, even if it occurs frequently, can no longer be learned. Hence, 1 (unknown) would be returned for every such input pattern. On the other hand, we envision a mechanism that recognized frequently present input patterns, while sporadic input patterns shall be classified as unknown. Therefore, we have changed the common ART neural network design to implement short-term memory that is based on an aging mechanism (see Section 3). Replacing sporadically matched prototypes is reasonable as frequently matched prototypes (normal input) are hardly affected by the aging mechanism. With this mechanism detection and learning can be continuously maintained.

3.3 System-Wide Anomaly Detection

In our basic system implementation, ART neural networks are implemented both at node level and at the base station. Thresholds are often used in related work. Therefore, we also evaluated simple threshold-based decisions. In these cases, a 1 (unexpected input vector) is reported if the sum of \mathbf{x} exceeds a predefined threshold T , i.e., if $\sum_{i=1}^{10} x_i > T$. Considering $\mathbf{x}_6 = [0, 0, 0, 33, 0, 0, 0, 0, 41, 13]$ above,

the sum is 87. Accordingly, if T is smaller than 87, a 1 is reported. Otherwise, no report is sent (the input is considered as expected or normal). The following four combinations are possible:

Local	System-wide
Fuzzy ART	binary ART
Fuzzy ART	Threshold-based
Threshold-based	binary ART
Threshold-based	Threshold-based

In the following evaluation the two crossed-out combinations are not considered: feeding local binary output to a threshold-based decider at the base station is not feasible. An example illustrates the corresponding problem. Consider two binary input vectors at the base station: $\mathbf{z}_1 = [1, 0, 0, 0, 0, 0, 0, 0, 0, 0]$ is an event that should be reported, while $\mathbf{z}_2 = [0, 1, 1, 1, 0, 0, 0, 0, 0, 0]$ is an event that should be filtered. This potential case is not manageable with a threshold-based decider at the base station, because $\sum_{i=1}^{10} z_{1,i} = 1$ (should be signaled), while $\sum_{i=1}^{10} z_{2,i} = 3$ (should be filtered), which conflicts with a threshold-based decision. Next, we consider a system with threshold-based decisions on node level and at the base station. It is important that $\sum_{i=1}^{10} x_i$ is reported to the base station instead of 1 if $\sum_{i=1}^{10} x_i > T$. Otherwise, the problem could be reduced to the previous one. This completely threshold-based approach cannot be parameterized such that the resulting anomaly detection system works reliably.

The designs of the remaining two systems are depicted in Figure 3. In every monitoring period of 20 s, every sensor node j signals a locally unknown event pattern either if the current time vector \mathbf{x}_j has not been recognized by the memory of node j (Fuzzy ART, see Figure 3(a)) or if $\sum_{i=1}^{10} x_i > T$ (threshold-based, see Figure 3(b)). Hence, the base station collects and processes space vectors $\mathbf{z} = \{z_1, \dots, z_j; j = 10\}$, where j represents the ID of the respective sensor

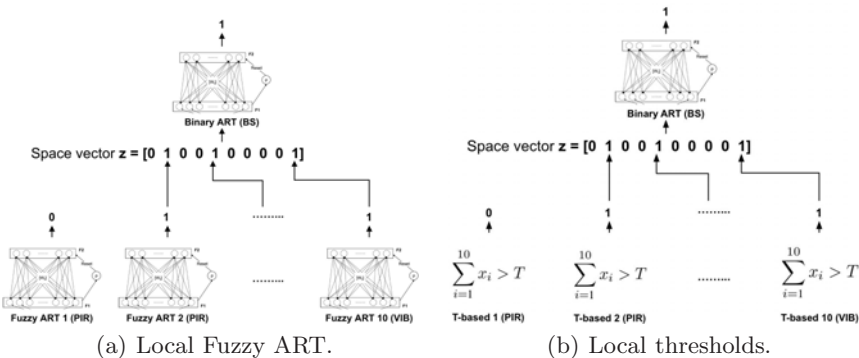


Fig. 3. System designs with local event recognition based on Fuzzy ART neural networks or thresholds T

node and $z \in \{0, 1\}$. Thus, if an unknown event pattern (signaled as 1) has been reported by a specific sensor node j , neuron j in F1 of the binary ART neural network is activated, i.e., 1 is fed to the corresponding neuron. In Figure 3 this means that 1 is fed to neurons 2, 5 and 10 of the binary ART neural network. If the base station did not receive a report from a sensor node, the base station assumes a recognized event pattern at the respective node. Hence, the corresponding neuron is fed with 0. This mechanism saves a considerable amount of reporting costs, because sensor nodes do not need to report known event patterns. The resulting input space vector \mathbf{z} in Figure 3 is $[0, 1, 0, 0, 1, 0, 0, 0, 0, 1]$. In Figure 3 this input vector is not known by the binary ART memory and a system-wide unknown event is reported, i.e., the output of the binary ART is 1.

Since single system-wide event reports are not sufficient to accurately signal anomalies (office intrusions), a significance test determining accumulations of system-wide event reports is provided. This test evaluates the frequency of anomalies over a certain time period.

```

Significance Test
Significance  $\Theta = 0$ ;
Age of last unknown event event_age = 0;
while true
  Calculate binary ART output  $\xi \in \{0, 1\}$ ;
  if  $\xi == 1$  // unknown event
    if event_age <  $T_{\max\_age}$ 
       $\Theta++$ ;
      event_age = 0;
    else
       $\Theta = 0$ ;
      event_age = 0;
    end
  else // known event
    event_age++;
  end
  if  $\Theta > T_{\text{significance}}$ 
    report anomaly;

```

The significance test is described in the pseudo-code above. The significance test determines an alarm, if in a certain time interval 5 system-wide unknown events ($\Theta > T_{\text{significance}}$, $T_{\text{significance}} = 4$) have been signaled by the binary ART neural network. The time difference between every subsequent pair of events must be smaller than 80 s ($T_{\max_age} = 5$, which implies $4 \cdot 20$ s).

Since five system-wide events are required, at maximum 320 s ($4 \cdot 80$ s) can elapse until an office intrusion is reported. On the other hand, the minimum delay is 80 s if 5 events are subsequently signaled ($4 \cdot 20$ s). The estimation of the maximum delay assumes that the office intrusion triggers unknown events. If this is (temporarily) not the case, the anomaly detection is delayed or disabled. The thresholds $T_{\text{significance}}$ and T_{\max_age} have been determined in simulations. Various thresholds have been evaluated. The used values have performed well.

4 Anomaly Detection Performance

The last part of this paper addresses the current deployment and its evaluation. As presented before, 10 sensing nodes have been deployed in an office offering two workplaces.

4.1 Office Occupancy Patterns

All experiments lasted between two and four days. Within these monitoring periods either normal office occupation or normal office occupation extended with specific hourly office access patterns were monitored. The specific patterns were either office searching performed by one person in an empty office room or hourly stress situations where multiple office staff were present, asking each other to look for missing items. The office access and occupation patterns are defined as follows:

- **Office searching:** The office is hourly searched for 2 - 5 minutes. The searching person enters the room and arbitrarily searches all different cabinets and drawers in the office. To avoid systematic search patterns, the office searching is performed by different persons. This pattern represents illegal access or abnormal behavior and is assumed to trigger alarms.
- **Stress situation:** In this office occupation pattern two to three persons are present in the office and are looking for some missing item(s). For example a document might be requested by an entering person and the two office personnel search and provide the requested information. Events of this kind last between 90 seconds and 3 minutes. This pattern imposes a high stress level on the system, but should not trigger alarms.
- **Normal office occupation:** Here, no restrictions have been defined. The office was just monitored for a given amount of time.

As discussed before, the system processes signals (PIR and vibration) in 20 s intervals on node level. Thus, every 20 s, local and global decisions are determined. If nothing happens in the office or the activation is recognized locally, no event reports are sent to the base station. Hence, the base station has input vectors of form $[0, 0, 0, 0, 0, 0, 0, 0, 0, 0]$. Such input vectors do not produce events at the base station. On the other hand, if many events occur during a certain time interval (office searching), many different unknown-event reports are sent to the base station. The resulting varying input vectors \mathbf{z} cannot all be known at the base station. Hence, the base station signals accumulations of events which leads to the system-wide alarm (anomaly).

4.2 Computation at Desktop PC

The system based on Fuzzy ART neural networks at the sensor nodes and on a binary ART neural network at the base station has been implemented and evaluated in the sensor network. On the other hand, determining the set of applicable

thresholds for the threshold-based local decider requires repeatable experiments. This cannot be provided by real-world experiments. Therefore, in addition to the running Fuzzy ART based system, every output of the sensor nodes was collected at a desktop PC, where the threshold-based anomaly detection system was implemented. Thus, the threshold-based decider could be evaluated and optimized offline with varying thresholds. To justify this implementation, the Fuzzy ART neural network based system was re-implemented at the base station, too. The results computed at the base station and at the desktop PC were equal. Hence, an offline analysis of the threshold-based decider is justified.

4.3 Detection Performance of the Anomaly Detectors

We investigated the number of false positives (false alarms) and false negatives (missing alarms) generated by the local Fuzzy ART and threshold-based (T-based) anomaly detectors. The threshold-based local anomaly detector has been evaluated with the complete range of applicable thresholds. Reports are generated if $\sum_{i=1}^{10} x_i > T$. Since analog input time vectors \mathbf{x} are processed and single sensor readings are in a range from 0 to 50, the resulting threshold T could be in the range from 0 to 500. High values are improbable because too many signals would be filtered. Comparatively low values of 16 and 17 for T have shown good prevention of false alarms, while no real alarms have been missed. All three office occupation scenarios were evaluated: normal office occupation, hourly office searching and hourly stress situations. The first experiment lasted for 48 hours including two working days. The two other kinds of experiments lasted for four days. In these four days, the respective occupation pattern (searching | stress) was performed eight times in eight hours. Accordingly, the latter two experiments provided 32 specific office occupation patterns (searching | stress) in each case.

Table 1 shows the anomaly detection performance of both anomaly detectors. The first important result is that, apart from one failure of the threshold-based system with $T = 16$, no false negatives were observed in all experiments. This system behavior is very important, because false negatives mean undetected intrusions. Of course, the presence of false negatives would question any anomaly detection and alarming system. On the other hand, some false positives, i.e., false

Table 1. False Positives (FP) and False Negatives (FN) of the anomaly detectors

	Fuzzy ART		T-based 16		T-based 17	
	FP	FN	FP	FN	FP	FN
Normal office occupancy (48h)	1	-	1	-	1	-
32 hourly office searchings (96h)	-	-	2	1	2	-
32 hourly stress situations (96h)	1	-	12	-	9	-
Total (240h)	2	-	15	1	12	-

alarms, could not be prevented. In particular the hourly stress level experiments generated false alarms, whereby the threshold-based system performed much worse. If no false positives can be tolerated, the system could be used to trigger a secondary high-resolution system (see also Section 3). In such an implementation the presence of false positives is less severe since only the secondary system is unnecessarily triggered.

The reporting of false alarms in the hourly stress level experiments is due to similarity of these experiments and office searching. The experiments have shown that Fuzzy ART neural networks are able to recognize and filter local anomalies, which leads to the system-wide prevention of false alarms. We conclude this section by highlighting that the local event recognition feature of the Fuzzy ART neural networks is in particular beneficial in presence of high stress level without intrusion. In this case, the system based on Fuzzy ART neural networks led to 1 false alarm in 32 stress situations, whereas the system based on local thresholds reported between 9 and 12 false alarms in the same situations. No false negatives (missing alarms) were encountered by both systems.

4.4 Message Load

Only normal office occupation was evaluated to assess communication costs. The other two experiments do not reflect normal behavior. In these intrusion or stress situations artificial anomalies are generated, which leads to temporarily increased communication load compared to normal daily office occupation. Signaling every local input vector \mathbf{x} with $\sum_{i=1}^{10} x_i > 0$ as event, i.e., threshold-based local decisions with $T = 0$, determines maximum possible communication costs. The costs of the respective anomaly detector have been computed in percentage of the maximum possible communication costs.

	Fuzzy ART	T-based 16	T-based 17
Monitoring without intrusion (48h)	11.9 %	7.0 %	5.9 %

The communication costs of the Fuzzy ART system and the T-based system are shown in percentage in Table 4.4. The local filtering of the Fuzzy ART neural networks leads to communication cost savings of up to 90%. Even better results can be achieved with the threshold-based filtering. However, the T-based system requires learning and training, i.e., the thresholds need to be determined. In contrast, Fuzzy ART neural networks are completely self-learning, i.e., only the memory size and the vigilance factor ρ need to be defined in advance. Furthermore, results in the last section have shown that the Fuzzy ART neural network approach prevents false alarms in case of stress situations.

Overall, without local filtering 10261 messages were reported to the base station during 48 hours. With the Fuzzy ART system this message load could be decreased to 1222 messages, which leads in average to 25 messages per hour. Considering the network size of 10 nodes and the faced office monitoring problem, a system-wide communication load of 25 messages per hour seems adequate.

4.5 Reporting and Triggering Delay

The reporting and triggering latency introduced by our anomaly detection system is investigated in this section. Section 3.3 has shown that the current implementation introduces a minimum reporting delay of 80 s and a maximum reporting delay of 320 s.

Table 2. Reporting Latencies [s] until alarm is reported

	Fuzzy ART		T-based 16		T-based 17	
	μ	σ	μ	σ	μ	σ
Hourly office searching (32h)	148	44	146	38	150	52

The effectively measured average reporting delays μ and their standard deviations σ are listed in Table 2. There is no significant difference between the different implementations. In average approximately 150 s were needed to detect and report office searching. The standard deviations vary slightly between 40 and 50 s. Detection latencies longer than 2 minutes seem rather long. However, this value depends on the duration of the monitoring period. The application of shorter monitoring periods could be evaluated. The monitoring period cannot be reduced arbitrarily due to the involved communication increase, though. The minimal reporting delay achievable is further restricted by the minimum amount of time required to identify abnormal behavior. We assume that reporting latencies of around 90 seconds might be achievable.

5 Conclusions and Future Work

In this paper a wireless sensor network for building monitoring has been proposed. The system detects and reports abnormal office occupancy. In contrast to conventional video surveillance systems, the system conceals the identity of the monitored office staff. Moreover, the deployed system is efficient and lightweight and produces much less data, decreasing administration and storage cost.

The system implements either a Fuzzy ART neural network or a simple threshold-based decider on a local level. A binary ART neural network is implemented at the base station. The proposed system with local Fuzzy ART decisions worked well in all experiments. In particular no false negatives were encountered, i.e., no cases of office searching were missed by the system. In normal office occupation with low stress level the threshold-based approach has performed similar to the Fuzzy ART neural network. Considering high stress levels, the Fuzzy ART neural network produces considerably fewer false positives than the threshold-based approach, though. This is due to the local memory maintained by the Fuzzy ART neural network that provides local recognition and filtering. It has been shown that communication costs could be cut by 90% with the Fuzzy ART-based system. The detection delay is currently approximately 2 minutes and 30 s, but could be further decreased by optimizing the monitoring cycles.

References

1. Lee, D.U., Kim, H., Tu, S., Rahimi, M., Estrin, D., Villasenor, J.D.: Energy-optimized image communication on resource-constrained sensor platforms. In: Proc. of the 6th international conference on Information processing in sensor networks (IPSN 2007), Cambridge, Massachusetts, USA, pp. 216–225 (2007)
2. Chen, M., Leung, V.C.M., Mao, S., Yuan, Y.: Directional geographical routing for real-time video communications in wireless sensor networks. *Computer Communications* 30(3368–3383), 17 (2007)
3. Xie, D., Yan, T., Ganesan, D., Hanson, A.: Design and implementation of a dual-camera wireless sensor network for object retrieval. In: Proc. of the 7th international conference on Information processing in sensor networks (IPSN 2008), Washington, DC, USA, pp. 469–480 (2008)
4. Luh, W., Kundur, D., Zourntos, T.: A novel distributed privacy paradigm for visual sensor networks based on sharing dynamical systems. *EURASIP Journal of Applied Signal Processing* 2007(1), 1–17 (2007)
5. Basharat, A., Catbas, N., Shah, M.: A framework for intelligent sensor network with video camera for structural health monitoring of bridges. In: Proc. of the Third IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOMW 2005), Washington, DC, USA, pp. 385–389 (2005)
6. Li, D., Wong, K.D., Hu, Y.H., Sayeed, A.M.: Detection, classification and tracking of targets. *IEEE Signal Processing Magazine* 19(2), 17–29 (2002)
7. Wang, T.Y., Han, Y.S., Varshney, P.K., Chen, P.N.: Distributed fault-tolerant classification in wireless sensor networks. *IEEE Journal on Selected Areas in Communications* 23(4), 724–734 (2005)
8. Wittenburg, G., Terloth, K., Villafuerte, F.L., Naumowicz, T., Ritter, H., Schiller, J.: Fence monitoring - experimental evaluation of a use case for wireless sensor networks. In: Langendoen, K.G., Voigt, T. (eds.) *EWSN 2007*. LNCS, vol. 4373, pp. 163–178. Springer, Heidelberg (2007)
9. Wälchli, M., Braun, T.: Event classification and filtering of false alarms in wireless sensor networks. In: Proc. of 3rd International Workshop on Intelligent Systems Techniques for Ad hoc and Wireless Sensor Networks (IST-AWSN 2008), Sydney, Australia (December 2008)
10. Kulakov, A., Davcev, D.: Intelligent wireless sensor networks using fuzzyart neural-networks. In: Proc. of the 12th IEEE Symposium on Computers and Communications (ISCC 2007), Aveiro, Portugal, July 2007, pp. 569–574 (2007)
11. Römer, K.: Discovery of frequent distributed event patterns in sensor networks. In: Verdone, R. (ed.) *EWSN 2008*. LNCS, vol. 4913, pp. 106–124. Springer, Heidelberg (2008)
12. Benbasat, A.Y., Paradiso, J.A.: A framework for the automated generation of power-efficient classifiers for embedded sensor nodes. In: Proc. of the 5th international conference on Embedded networked sensor systems (SenSys 2007), Sydney, Australia, pp. 219–232 (2007)
13. Gu, L., Jia, D., Vicaire, P., Yan, T., Luo, L., Tirumala, A., Cao, Q., He, T., Stankovic, J.A., Abdelzaher, T., Krogh, B.H.: Lightweight detection and classification for wireless sensor networks in realistic environments. In: Proc. of the 3rd international conference on Embedded networked sensor systems (SenSys 2005), San Diego, California, USA, November 2005, pp. 205–217 (2005)
14. Cai, J., Ee, D., Pham, B., Roe, P., Zhang, J.: Sensor network for the monitoring of ecosystem: Bird species recognition. In: Proc. of the 3rd International Con-

- ference on Intelligent Sensors, Sensor Networks and Information (ISSNIP 2007), Melbourne, Australia, pp. 293–298 (2007)
15. Powell, G., Marshall, D., Smets, P., Ristic, B., Maskell, S.: Joint tracking and classification of airbourne objects using particle filters and the continuous transferable belief model. In: Proc. of the 2006 9th International Conference on Information Fusion (Fusion 2006), Florence, Italy, July 2006, pp. 1–8 (2006)
 16. Wang, X.R., Lizier, J.T., Obst, O., Prokopenko, M., Wang, P.: Spatiotemporal anomaly detection in gas monitoring sensor networks. In: Verdone, R. (ed.) EWSN 2008. LNCS, vol. 4913, pp. 90–105. Springer, Heidelberg (2008)
 17. Patcha, A., Park, J.M.: An overview of anomaly detection techniques: Existing solutions and latest technological trends. *Computer Networks* 51(12), 3448–3470 (2007)
 18. Roman, R., Zhou, J., Lopez, J.: Applying intrusion detection systems to wireless sensor networks. In: Proc. of the 3rd IEEE Consumer Communications and Networking Conference (CCNC 2006), Las Vegas, Nevada, USA, January 2006, vol. 1, pp. 640–644 (2006)
 19. Krontiris, I., Dimitriou, T., Giannetsos, T., Mpasoukos, M.: Intrusion Detection of Sinkhole Attacks in Wireless Sensor Networks. In: *Algorithmic Aspects of Wireless Sensor Networks*, pp. 150–161. Springer, Heidelberg (2008)
 20. Li, Y.Y., Parker, L.: Intruder detection using a wireless sensor network with an intelligent mobile robot response. In: Proc. of the IEEE SoutheastCon 2008, Huntsville, Alabama, USA, pp. 37–42 (2008)
 21. Dasgupta, D., Forrest, S.: Novelty detection in time series data using ideas from immunology. In: Proc. of The Fifth International Conference on Intelligent Systems (IS 1996), Reno, Nevada, USA (1996)
 22. Mazhar, N., Farooq, M.: A sense of danger: dendritic cells inspired artificial immune system for manet security. In: Proc. of the 10th annual conference on Genetic and evolutionary computation (GECCO 2008), Atlanta, GA, USA, pp. 63–70 (2008)
 23. Burwick, T., Joublin, F.: Optimal algorithmic complexity of fuzzy art. *Neural Processing Letters* 7(1), 37–41 (1998)
 24. Scatterweb: The self-organizing wireless communication platform (October 2007)
 25. Carpenter, G., Grossberg, S.: *Adaptive Resonance Theory*. Bradford Books. MIT Press, Cambridge (2002)