Low Overhead Assignment of Symbolic Coordinates in Sensor Networks

Matthias Gauger · Pedro José Marrón · Daniel Kauker · Kurt Rothermel

Abstract Approximate information on the location of nodes in a sensor network is essential to many types of sensor network applications and algorithms. In many cases, using symbolic coordinates is an attractive alternative to the use of geographic coordinates due to lower costs and lower requirements on the available location information during coordinate assignment. In this paper, we investigate different possible methods of assigning symbolic coordinates to sensor nodes. We present a method based on broadcasting coordinate messages and filtering using sensor events. We show in the evaluation that this method allows a reliable assignment of symbolic coordinates while only generating a low overhead.

1 Introduction

Information on the context of nodes in a sensor network is essential to many types of sensor network applications. Examples of such context knowledge are the positions of nodes, their neighborhood or the external conditions of a node's surroundings. Among the different types of context, location information plays an especially important role in sensor networks as it is required if sensor readings are to be associated with the area they were recorded in.

Acquiring position information of sensor nodes in the form of geographic coordinates with an acceptable precision is a very difficult and often costly operation. This is especially true for indoor scenarios where localization technologies like GPS receivers do not work well. One possible alternative to determining geographic coordinates is to assign symbolic coordinates to nodes. Instead of describing positions in the form of a coordinate tuple, a symbolic coordinate represents areas of different shapes and sizes in the form of an abstract symbol. All sensor nodes in an area have the same symbolic coordinate. Examples of such symbolic coordinates are room numbers in a building or street addresses.

In several types of sensor network applications it is possible to use symbolic coordinates instead of geographic coordinates. One exemplary field is the retrieval of sensor data from specific areas of a sensor network. In many cases, symbolic coordinates directly represent the semantics of a location, for example, when a symbolic coordinate is associated with each room of a building. This allows to implement data retrieval operations very easily without having to map from coordinates to areas first. Another possible application of symbolic coordinates is using them for cost-effective many-to-many routing of messages in sensor networks which we have demonstrated in previous work [5]. We specify symbolic source routes from the sender node to the destination area and later translate these routes on the node level into specific routing decisions.

The applications of the methods presented in this paper do not have to be limited to the assignment of symbolic coordinates. It is also possible to use them for node clustering (i.e., all nodes in a room form a cluster). In such room-level clusters the nodes typically generate related information which can be aggregated with less information loss than in arbitrary clusters.

We argue that assigning symbolic coordinates in sensor networks is usually much easier than assigning

Matthias Gauger · Pedro José Marrón · Daniel Kauker
Universität Bonn, Bonn, Germany and Fraunhofer IAIS, St. Augustin, Germany
E-mail: {gauger, pjmarron, kauker}@cs.uni-bonn.de

Kurt Rothermel
IPVS, Universität Stuttgart, Stuttgart, Germany
E-mail: rothermel@ipvs.uni-stuttgart.de

Published in Telecommunication Systems, Springer Netherlands, Volume 40, Numbers 3-4 / April 2009, pp 117-128

geographic coordinates and is possible with reasonable effort. In this paper, we discuss different approaches to this assignment for indoor scenarios and present one solution that achieves a very low error rate while only generating a small overhead and requiring no prior knowledge on the sensor network topology. The basic idea is to let an administrator broadcast symbolic coordinates in the different rooms a sensor network is deployed in and let the nodes use sensor information to filter out broadcast messages wrongly received from neighboring rooms.

The rest of this paper is organized as follows. The following section briefly reviews important related work. After that, we present routing based on symbolic coordinates as an important example application motivating this work in Section 3. We present our three approaches to the assignment of symbolic coordinates in Section 4. In Section 5 we describe relevant details of our implementation before Section 6 provides an in-depth evaluation of our approaches. Section 7 concludes the paper and discusses some future work.

2 Related Work

Our work is related to the large set of approaches in the area of node localization which is one of the fundamental research problems in wireless sensor networks. Most localization approaches have in common that they require a set of anchor nodes with known positions. Then different techniques and algorithms are used to determine the positions of the other nodes (see for example [8]). One important factor is how distances between nodes are determined. There are range-free solutions that only consider the hop-count (e.g., [10]), distance estimations based on the received signal strength (e.g., [9]) and solutions requiring special ranging devices (e.g., [11]). Elnahrawy, Li and Martin [4] discuss the fundamental limits of localization techniques based on signal strength when used in indoor scenarios.

The Spotlight localization system [12] is particularly related to our approach in that it also uses sensor events for the localization of nodes. A helicopter (the Spotlight device) which knows its own position flies over the sensor network and generates light events at certain points of time. The sensor nodes report when they detect events back to the helicopter, which is then able to compute the geographic coordinates of the nodes. However, as the authors aim to calculate geographic coordinates, the required calculations are rather complex and a precise time synchronization of nodes is required.

StarDust [13] passively localizes sensor nodes with the help of reflected light. Each sensor node is equipped with a retroreflector which reflects light coming from an aerial vehicle. This vehicle records an image of the deployment area with a digital camera showing the light reflections and image processing techniques are used to identify the locations of the nodes. A mapping of nodes to the set of determined positions is found with the help of a relaxation algorithm based on information like the neighborhood relationship among nodes.

Similar to our assignment of symbolic coordinates, Corke, Peterson and Rus [3] assign geographic coordinates to sensor nodes using radio communication. In their scenario, a robot helicopter equipped with a GPS receiver flies over the network area and periodically broadcasts beacon messages containing its current geographic coordinate. The sensor nodes on the ground typically receive multiple such beacons and need to estimate their position based on this information. The authors propose different methods for this calculation, including taking the mean of the positions or the signal strength weighted mean of the positions. A constraint-based method performed best in their experiments.

There has been active work on symbolic coordinates, their applications and underlying location models mostly in the area of pervasive computing (e.g., [2], [7]). Becker and Dürr [1] give a comprehensive overview of different geometric and symbolic location models from the perspective of pervasive computing and compare their suitability for different types of queries. In our work we assume that such a symbolic location model is available so that an administrator is able to assign symbolic coordinates to individual rooms.

3 Routing Based on Symbolic Coordinates

In this section we describe our approach for routing based on symbolic coordinates [5] in a little more detail as one exemplary application of symbolic coordinates that illustrates the benefit of having symbolic coordinate information available in wireless sensor networks.

The problem to be solved is the routing of messages between mobile devices and resource-constrained wireless sensor networks. Typical sensor network routing solutions are too restricted to support multiple clients and client mobility (e.g., all communication is done through a central base station) while routing approaches from the wireless ad-hoc networks domain are not well-suited for resource-poor sensor nodes (for example requiring the maintenance of global routing tables).

The basic idea of our approach is to perform a *symbolic source routing* in sensor networks by letting a mobile client node specify a symbolic source route from its current position to the destination area. The static sensor nodes can take advantage of this source route to

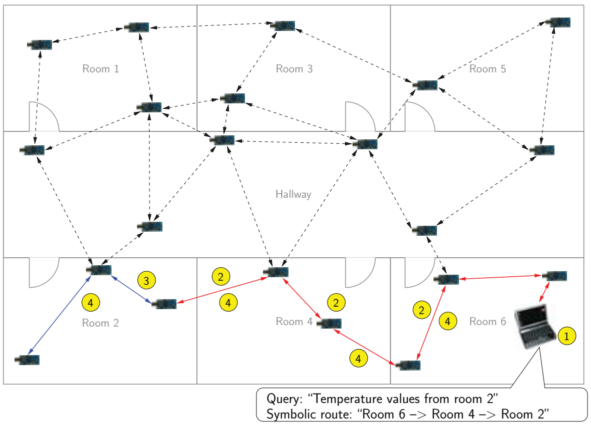

Fig. 1 Routing example

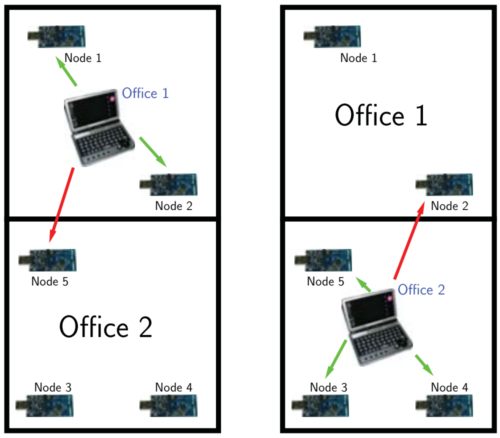

(a) (b)

Fig. 2 Example scenario

make node-to-node routing decisions based on purely local routing information. Fig. 1 shows a simple example of a mobile client device querying the sensor network in an office scenario. In addition to the query itself ("Temperature values from room 2") the client also provides the symbolic route ("Room 6 - Room 4 - Room 2") as part of the query message. The sensor nodes have learned about their local neighborhood using a beaconing mechanism and maintain routing information to neighboring symbolic coordinates based on this neighborhood information. In the example, the sensor nodes in "Room 6" know how to forward messages in order to reach "Room 4" and nodes in "Room 4" know how to reach the neighboring area "Room 2". Note, however, that the size of these routing tables and the effort for maintaining them is limited as they only contain information on the local symbolic neighborhood.

The main advantage of splitting the routing task between mobile nodes and static sensor nodes as described above is that the mobile client nodes do not need to manage a detailed and up-to-date view of the sensor network topology. Such a global view on the topology would be hard to accomplish given the large size of typical sensor networks and the relative instability of the network links. At the same time, the sensor nodes can correctly forward messages using purely local information which effectively limits the costs for the resource-poor devices.

With the help of the mechanisms described in this paper it is possible to assign symbolic coordinates to sensor nodes with low overhead making this routing approach an attractive alternative to other routing solutions, like for example a routing based on geographic coordinates.

4 Assignment of Symbolic Coordinates

In this section we introduce and describe our three different approaches to assign symbolic coordinates in sensor networks. We start by describing the properties of our target scenarios in more detail.

Our target system consists of a set of sensor nodes that are distributed to different rooms in a building. Each room is uniquely identified by a symbolic coordinate, which can, for example, correspond to the respective room number. We assume that each sensor node is preprogammed with a node identifier (node ID) and that this node ID is unique in the sensor network. However, we do not assume that there is any relation between this node ID and the location of the node in the network. In the beginning, the sensor nodes do not have any information about the symbolic coordinate of the room they are located in.

After the deployment of the sensor nodes there is one person – who we call administrator – responsible for assigning the correct symbolic coordinates to the nodes in the network. The administrator has a mobile client device that is able to directly communicate with the nodes of the sensor network. The client device can be used to send so-called coordinate messages to the nodes.

Fig. 2 illustrates our system model with a simple example scenario with five sensor nodes distributed in two rooms. The administrator first visits the room "Office 1" (see Fig. 2 (a)) and then the room "Office 2" (see Fig. 2 (b)).

4.1 Individual Assignment of Symbolic Coordinates

The most basic way of assigning symbolic coordinates to nodes is to assign the respective coordinate to each node individually. In some scenarios it might be possible to directly encode the symbolic coordinate in the program code of the node like it is typically done with the node ID. However, this limits the flexibility for the placement of nodes. Moreover, we generally expect the sensor nodes to be delivered to users preprogrammed with an application software when the application field lies outside of typical research settings.

To assign symbolic coordinates to sensor nodes after the deployment of the network, the administrator needs to send the correct symbolic coordinate in a coordinate message to each individual node using the node's wireless interface. Upon receiving such a message, the node stores the new symbolic coordinate and uses it from this point on.

The clear advantage of an individual assignment of symbolic coordinates to nodes is that it avoids ambiguity: The administrator has complete control over which nodes receive which symbolic coordinate and he is able to ensure the correctness of the assignment process. However, there are also a number of clear disadvantages. The assignment of coordinates requires individual communication with each sensor node in the network. Consequently, the required time and effort (and also the message complexity) grow with the number of nodes in the network. More important, it is necessary that the administrator has an up-to-date knowledge on the distribution of nodes to the individual rooms in the building (including information on which node identifier belongs to which node) which might require a time-consuming visual inspection of the complete sensor network.

4.2 Assignment of Symbolic Coordinates by Broadcast

Our second approach, the assignment of symbolic coordinates by broadcast, aims to address the disadvantages of having to separately assign a symbolic coordinate to each individual node in the network. Instead, the goal is to distribute the coordinate information to all nodes in an area in a single step.

The network administrator needs to visit the different rooms covered by nodes of the network and has to send out a message containing the current symbolic coordinate information in each room. The message is sent by broadcast so that all nodes in the one-hop neighborhood of the client device receive the information. Upon receiving such a coordinate message, a sensor node overwrites its symbolic coordinate information with the newly received data.

One advantage of sending out coordinate information by broadcast is the lower overhead both for the administrator and in terms of the number of messages as the configuration has to be done only once per room instead of once per sensor node. Moreover, no information is required on the position of individual sensor nodes since the coordinate information does not have to be addressed to specific nodes. The only information that must be available is the symbolic coordinate of the room the broadcast message is sent in.

It is typically desirable that sensor networks are connected across area boundaries like walls between rooms to provide for communication between different network parts. Therefore, the main challenge in the assignment of symbolic coordinates by broadcast is nodes that receive coordinate messages from neighboring rooms. Depending on the sequence of messages sent, these nodes might overwrite the correct coordinate information with data belonging to a neighboring room. The problem is illustrated in the example in Fig. 2 where a broadcast sent in "Office 1" also reaches node 5 in "Office 2" and a broadcast sent in "Office 2" also reaches node 2 in "Office 1".

The basic approach to address this challenge is to control the signal strength of the messages sent by the client device. Ideally, the strength of the signal should still allow the message to reach all nodes inside the current room and none of the nodes in neighboring rooms. However, it is difficult or even impossible to find this balance especially if two nodes in adjacent rooms are located very close to each other and the attenuation of the signal by the wall between them is small.

We have developed two extensions for the assignment of symbolic coordinates by broadcast. The first extension allows sensor nodes to store and manage multiple symbolic coordinates at the same time so that different coordinate messages do not overwrite each other's information. If a node has received multiple symbolic coordinates, then it lies in the border area of the rooms represented by these coordinates. How these multiple coordinates of a node are used is application dependent. For example, a node might participate in the operations of all areas it supposedly is a member of.

The second extension to the assignment by broadcast analyzes the signal strength of the different coordinate messages received at a node, called Received Signal Strength Indication (RSSI) value, and uses this information to assign a coordinate to the node. It uses the symbolic coordinate from the message received with the highest RSSI value reflecting the assumption that a coordinate message sent from the same room should be received with a higher signal strength than coordinate messages from different rooms. However, this only works reliably if the signal attenuation of the walls between rooms is significant. This can be problematic because RSSI – despite its strong limitations in indoor scenarios [4] – is more of an indicator for geographic distances among nodes than an aid for the localization of nodes to different areas.

4.3 Assignment of Symbolic Coordinates by Assisted Broadcast

Both approaches to the assignment of symbolic coordinates presented so far have disadvantages. The first approach is quite intricate and requires detailed knowledge on the nodes' positions in the network whereas the second approach is susceptible to errors due to the propagation of coordinate messages across room boundaries. In the following, we present a third approach, the assignment by assisted broadcast, which avoids these problems.

Like for the assignment of coordinates by broadcast described above, the client device broadcasts coordinate messages. It is again the task of the administrator to set the symbolic coordinate sent out by the client in a way so that it corresponds to the current location of the administrator. The important difference lies in the handling of a received coordinate message by the sensor nodes. Instead of directly assigning a symbolic coordinate received in a coordinate message, a node first checks whether the new coordinate is confirmed by a sensor stimulus following the message.

Directly after the administrator has sent a coordinate message (within a pre-defined time interval of a few seconds) he needs to trigger an event that can be detected by one of the sensors of the nodes. The goal is to distinguish nodes that are located in the same room as the administrator and that both receive the coordinate message and detect the following sensor event from nodes in neighboring areas that only receive the coordinate message but are not affected by the sensor event triggered by the administrator.

Different types of sensors and different types of sensor stimuli can be used to confirm coordinate messages. The main preconditions are that the sensor stimuli can be easily triggered by the administrator and that the resulting sensor event can be detected in the whole area but not in neighboring areas. We work with two types of sensor events, namely light level changes and acoustic events. In indoor scenarios, the light level in rooms can easily be influenced by turning the artificial lighting on or off. Acoustic events can be generated and detected using special devices attached to the sensor nodes as described in the implementation section.

Which behavior constitutes an event and how such an event is detected depends on the type of sensor and the monitored characteristic of the environment. For light sensors we define an event as a significant change of the recorded luminance level within a limited time period. Fig. 3 shows the pseudo code for checking whether a light event has been detected and the received symbolic coordinate should be assigned. Directly after having received a new coordinate message, the node records its current sensor value and starts a timer. When the timer fires, it records its sensor value again and calculates the absolute difference of the two sensor values. If this value is above a specified event threshold, the sensor node detects an event and assigns the new symbolic coordinate.

We define events differently for acoustic signals: A sensor node detects an event when it detects a sound of a certain frequency. For the coordinate assignment by assisted broadcast this means that the symbolic coordinate is assigned to a node if it detects a pre-defined sound within a limited time period after it has received the coordinate message.

The method described above relies on two important assumptions. First, it assumes that it is possible for the administrator to change the external conditions in a way that allows all sensor nodes in the room to detect these changes as an event. Second, it also assumes that similar changes to the external conditions in the neighboring rooms are unlikely to happen at the same time without explicit intervention by an external party.

The main advantage of the assignment by assisted broadcast is that the additional sensor stimulus triggered by the administrator prevents ambiguities in the assignment of the symbolic coordinates. This allows to assign the coordinate to a specific set of nodes without having to address each of the nodes in this set individually. However, this comes at the cost of the additional effort required for generating the external sensor stimulus. Moreover, it only works for actual sensor nodes that possess the sensor chip required for detecting the event. The first two approaches also work on other devices that are part of the sensor network, for example gateway nodes without any sensing functionality.

5 Implementation

We have implemented our three approaches for the assignment of symbolic coordinates based on the TinyOS 2.0 operating system with support for different sensor node platforms (TelosB, MICAz, MICA2). For most experiments, we used TelosB sensor nodes which provide

```
int ownSymbolicCoordinate;
int candidateSymbolicCoordinate;

event receivedCoordinateMessage(int newSymbolicCoordinate) {
  candidateSymbolicCoordinate = newSymbolicCoordinate;
  initialSensorValue = getSensorValue();
  startTimer(eventDetectionTimerLength);
}

event timerFired() {
  int finalSensorValue = getSensorValue();
  if (abs(currentSensorValue - finalSensorValue) > eventThreshold) {
    ownSymbolicCoordinate = candidateSymbolicCoordinate;
  }
}
```

Fig. 3 Checking a new symbolic coordinate for applicability

sensor chips for measuring temperature, humidity as well as two light sensor chips that capture the photosynthetically active radiation (PAR) and the total solar radiation (TSR) respectively. For the evaluation of the assignment by assisted broadcast based on light events we used the TSR light sensors.

For our experiments with acoustic events we used MICA2 and MICAz sensor nodes in combination with MTS300 sensor boards from Crossbow. In addition to a light and a temperature sensor, the MTS300 contains a sounder element and a microphone. The sounder is able to emit sounds with a frequency of 4kHz. The microphone is connected to a hardware tone decoder that reports an event when a sound in the 4kHz frequency range is recorded by the microphone. This way, the administrator is able to trigger an acoustic event using the sounder element of an MTS300 sensor board while the sensor nodes use the microphone and the tone decoder to detect such events.

The client application is implemented in C++ based on the Qt software toolkit. It runs on Linux PDAs from Sharp (Sharp Zaurus SL-3200) that communicate with the sensor network using a TelosB sensor node connected to the PDA over USB as a bridge node. In the future, we expect the availability of devices that are able to directly communicate with sensor nodes using a communication standard like IEEE 802.15.4.

The client application supports all three approaches presented in this paper and both light and acoustic events. Thus, symbolic coordinates can be sent to individual nodes or broadcasted with or without advertising a following sensor stimulus. Note that acoustic events can be directly triggered by the client application by activating the sounder on the sensor node connected to the PDA after the coordinate message has been sent whereas triggering a light event requires explicit activity by the administrator. This is an important advantage of the solution using acoustic events.

The output power of the CC2420 radio chips used by the TelosB and the MICAz sensor nodes is programmable. The possible range of values starts with a minimum output power of -25 dBm and goes up to a maximum output power of 0 dBm. The client application allows to select any of the valid output power values.

6 Evaluation

In this section, we present an evaluation of our different approaches to the assignment of symbolic coordinates in sensor networks. However, we do not discuss the first approach in detail since the manual assignment of coordinates to individual nodes should work in all cases as long as no packets are lost due to the unreliability of the wireless channel.

For evaluating our approach we deployed 14 sensor nodes in 7 different rooms of the computer science building at the Universität Stuttgart. Fig. 4 shows the floor plan and outlines the location of the nodes in the rooms. Note that we tried to create a somewhat irregular distribution of nodes with different distances among nodes in different rooms.

In a first set of experiments we investigated how reliable the assignment of symbolic coordinates using broadcast works. For this purpose, we sent out coordinate messages in each room with different output power levels and collected information on which nodes received the coordinate message with which RSSI value. We repeated each experiment five times for each signal strength level of the client device.

As a first result, these experiments showed that effectively limiting the dissemination of coordinate messages to a single room is hardly possible even when using the minimum transmission power of the TelosB sensor nodes. To illustrate this, Fig. 5 (a) shows the

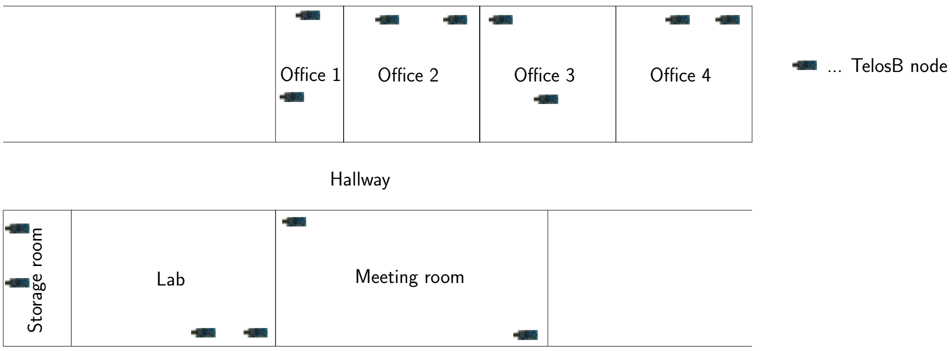

Fig. 4 Floor plan of the experiment area

Table 1 Average performance for different sender signal strengths

	Output power -25 dBm	Output power -14 dBm
Max. # of nodes assigned correctly	13	13
Min. # of nodes assigned correctly	11	10
Avg. % of nodes assigned correctly	88.57	77.14

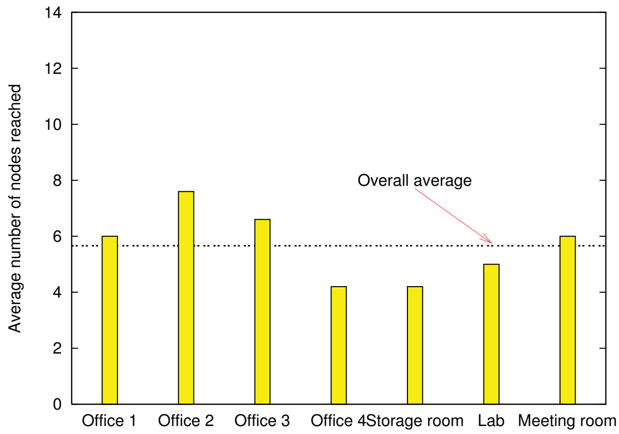

(a) Output power -25 dBm

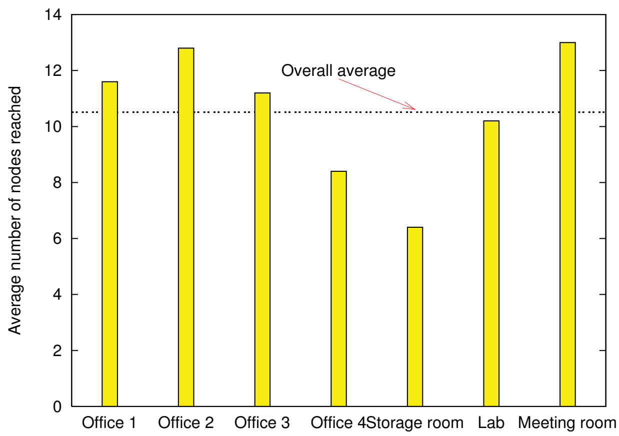

(b) Output power -14 dBm

Fig. 5 Average number of nodes reached by coordinate message broadcasts

average number of nodes reached from each room with the minimum output power of approximately -25 dBm and Fig. 5 (b) shows the same analysis for an output power of approximately -14 dBm. Even for the minimum output power, the overall average of 5.66 nodes reached from each area is much higher than the two nodes actually located in each area. On average, each node received coordinate messages from 2.83 different areas (5.26 for an output power of -14 dBm) with a maximum of 5 (7 for an output power of -14 dBm).

To deal with multiple coordinate messages received from different rooms we proposed to consider the RSSI values of the received messages and assign the symbolic coordinate from the message with the highest RSSI value. Using this extension results in a promising performance of the symbolic coordinate assignment using broadcast: Table 1 shows the maximum and the minimum number of nodes with correctly assigned symbolic coordinates for the experiments with signal output powers of -25 dBm and -14 dBm as well as the average percentage of nodes assigned correctly.

Coordinate assignments with the smaller transmission power level clearly outperform the assignment with a higher transmission power level and only assign the wrong coordinate to between one and three nodes. An explanation for this can be found looking at the RSSI values: On average, the RSSI values of messages sent in the same room are only 6% (8% for -25 dBm) larger than the RSSI values of messages sent from different rooms. While the ratio between inside and outside RSSI values is a little higher for -14 dBm than for -25 dBm this cannot compensate the much higher number of coordinate messages each node receives. Due to the inherent variations of the RSSI values of received messages, a larger number of coordinate messages received from neighboring rooms increases the probability that the

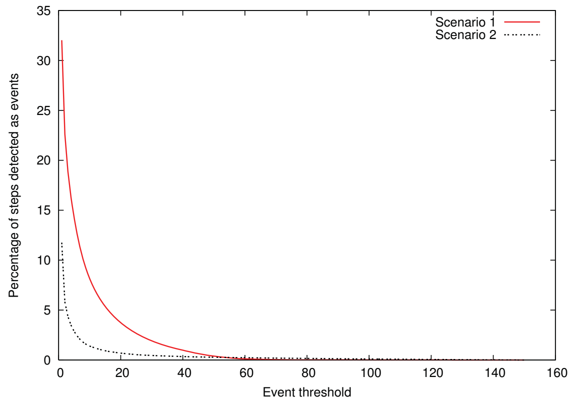

Fig. 6 Percentage of events detected for different event thresholds

RSSI value of one of these messages is larger than the RSSI value of the message received from the own room.

Overall, while the results of our second approach together with an analysis of the RSSI values produced good results in our experiments, our analysis also made clear that RSSI is a fragile criterion that is not able to produce 100% reliable results.

6.1 Assisted Broadcast Using Light Events

Based on the results of the described experiments we next investigated the assignment of symbolic coordinates by assisted broadcast using light events. We used the raw sensor data read from the analog-to-digital converter for our experiments without converting the data to a measuring unit like lux.

First, we wanted to investigate how often TSR light events typically occur when they are not explicitly triggered by the user. For this purpose, we collected sensor data in two indoor scenarios. We distributed 12 nodes to 4 different rooms in both scenarios and collected the value of the TSR light sensor every 10 seconds over multiple days with each sensor node. In the analysis, we evaluated how often an event is detected for two consecutive measurements when varying the event detection threshold. Fig. 6 shows the result of this analysis for both scenarios.

The results show that for very small event thresholds a considerable percentage of measurement pairs triggers events. However, with an event threshold of 10 only in 8.1% (1.4%) and with an event threshold of 20 only in 3.7% (0.7%) of the cases an event has been detected. This indicates a quite small probability of unintended events occurring during the coordinate assignment using these thresholds especially since the event detection time period (10 seconds) was selected quite large in this case. Since changes to the level of illumination of the sensor nodes during the daytime are the main sources of events, the considerable differences between scenario 1 and scenario 2 can be explained by the fact that more nodes in scenario 1 were exposed to sunlight than in scenario 2.

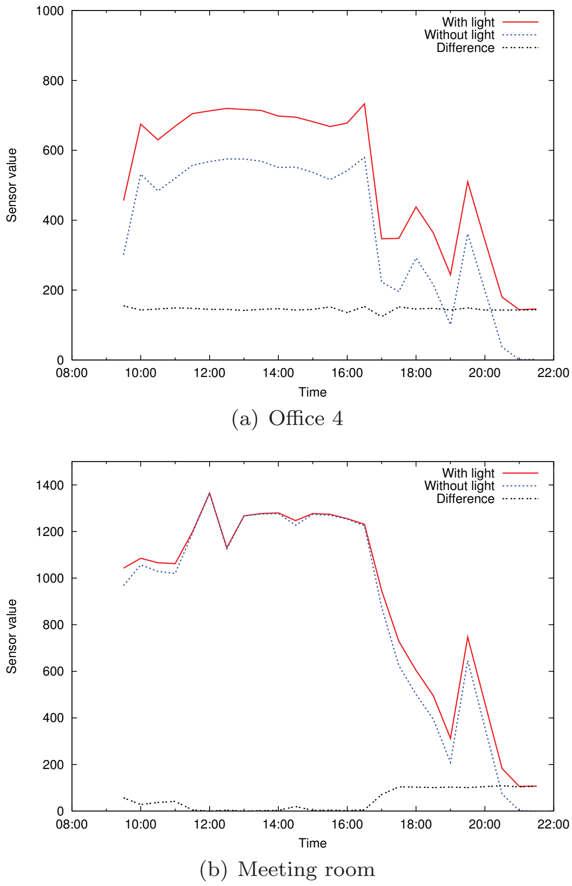

Fig. 7 Light sensor data with and without room light

Besides unintentionally detected events a second potential issue is how to trigger light events when the room is already brightly lit by sunlight coming through the window. To investigate how big of a problem this is we recorded TSR light sensor values in two of our rooms for one day. One recording was done in office 4 which has its windows to the north. To explore more extreme conditions we placed the sensor for the second recording directly behind one of the south-bound windows of the meeting room. Values were recorded every 30 minutes on a sunny day both with the light turned on and with the light turned off. Fig. 7 shows the resulting graphs for both rooms.

The difference between the sensor values recorded with and without light is relatively stable over the day in office 4 (Fig. 7 (a)). With difference values in the range of 140, a reliable detection of light events is possible irrespective of the time of day the coordinate assignment is done. The situation is different for the sensor at the window of the meeting room (Fig. 7 (b)). Here, the influence of the artificial light is considerably lower and the difference between the sensor values shrinks down to values around 0 at noon. Obviously, the artificial light in the room cannot add to the recorded sensor value anymore once a certain light level is reached in the environment. Therefore, the time period of the coordinate assignment must be selected carefully if nodes are placed at especially exposed locations. However, our experiences also show that for normal conditions event detection is possible all day if the event threshold is set to a reasonable value.

The last step of our evaluation is to investigate the results of using the third approach with different event detection thresholds. We performed a set of experiments choosing different event thresholds at different times of the day. In all of these experiments the client application sent the coordinate messages with the minimum possible transmission power and we used an event detection timer length of 4 seconds. Table 2 shows an overview of the results. Note that we deliberately ignored the RSSI values for this evaluation to emphasize the benefit of using events. However, the RSSI value could be used as an additional criterion.

Table 2 Average success rates for different event thresholds

Event detection threshold	Average percentage of nodes assigned correctly
5	53.57
10	100.0
20	100.0
100	78.57

If the event threshold is set to a too small value, then events can occur without being explicitly triggered by a user simply due to the normal variations of the sensor values over time. The consequence of this are so-called false positives during the assignment of symbolic coordinates – nodes that assign a coordinate in reaction to a coordinate message without lying in the room where the event has actually been triggered. In our experiments we could observe this for an event threshold of 5 (see Table 2). The results of experiments performed in the evening or at night with this threshold lay above the average but even then the artificial light oscillated enough to generate some false positives.

To explore the other end of the spectrum, i.e., a high event detection threshold, we performed experiments with a threshold value of 100. As expected, the high event detection threshold reliably prevented the occurrence of any false positives. However, some of the intended recipients also did not detect an event and consequently did not assign the symbolic coordinate resulting in a success rate well below 100%.

With an event threshold of 10 or 20 all of our experiments assigned the correct symbolic coordinates to all nodes in the network irrespective of the time of day the experiment was performed. Therefore, selecting an event threshold in this range provides for a reliable assignment of symbolic coordinates to sensor nodes.

6.2 Assisted Broadcast Using Acoustic Events

In addition to our experiments with the assignment of symbolic coordinates based on the detection of light events, we also experimented on a smaller scale with the assignment by assisted broadcast using acoustic events. The main purpose of these experiments was to determine whether acoustic events are a viable and useful alternative to light events, for example because they are not susceptible to the influence of daylight.

The first insight of our experiments with acoustic events on the MTS300 sensor boards was that false positives – acoustic events detected by the tone detector without a signal sent by a sounder – are a much more critical problem than for light events. The simple tone detection circuit on the MTS300 sensor boards does not react exclusively to the sounds emitted by the MTS300 sounder element but can also be triggered by a variety of other sounds. For example, in many cases the simple knocking on a table, a lively conversation or the sound of typing on a keyboard was able to trigger an event at nearby sensor nodes.

To evaluate the severity of this problem we recorded acoustic events in experiments in a set of different situations: We collected data in an office environment, in a room with a running radio and a very quiet room without any activity. We used the measurements from the office environment as an example of typical everyday activity to be found in many indoor scenarios where sensor nodes are likely to be deployed. The measurements from the room with a running radio represent the extreme case of a particularly noisy environment whereas the quiet room represents optimal conditions to be found during the coordinate assignment. In each of these experiments, a MICA2 sensor node continuously listened for sound events for a time period of one hour and incremented its counter for each occurrence of an event detected during this time.

Table 3 shows the results of the false positive experiments. Clearly, in noisy environments like our room with a running radio, the number of false positive events detected is unacceptably high with an event occurring every 0.59 seconds on average. While the number is significantly lower for the quieter office environment, a false positive event every 10.71 seconds still poses a high risk of errors during the assignment of symbolic coordinates using acoustic events. Note that for the experiment with the running radio, the recording sensor node was placed near the loudspeaker of the radio and the radio was operated at a relatively high volume. In another experiment where we placed the sensor node in a different part of the room than the radio, the results for the number of false positive events resembled the results of the standard office scenario. As expected, in a quiet room the number of false positives remains in the range of zero. However, it is probably difficult to guarantee such conditions during the whole assignment of symbolic coordinates in a building.

Table 3 False positive analysis acoustic events

Scenario	**Number of false positive events in 1 hour**
Office environment	336
Radio running	6068
Quiet room	2

Even though the sensors are only sampled in reaction to a received coordinate message, we still considered the rate of false positive events as too high for a reliable assignment of symbolic coordinates. Our solution to this problem is to extend the event detection process and to listen for a sequence of sound events instead of a single event. The underlying assumption is that the probability of the random occurrence of a sequence of sounds (or one long-lasting sound) detectable as events is smaller than for the occurrence of a single sound event.

The extended event detection works as follows: After receiving a coordinate assignment message, a sensor node starts a timer that fires after x milliseconds. It checks whether an acoustic event occurred during this time interval. The timer is restarted a maximum number of y times. If more than z acoustic events are detected during these y time intervals, then the sequence is marked as detected and the symbolic coordinate is assigned. As an example, a node waits for 4 time periods with a length of 500 milliseconds each and is required to detect an acoustic event in 3 of these time periods to actually assign the symbolic coordinate it received. Only requiring the detection of z out of y possible sensor events reflects our experiences during preliminary experiments where none of the sensor nodes was able to detect all sensor events triggered by a nearby sounder.

Table 4 shows the number of false positive events recorded in the office environment and in the scenario with a running radio for three exemplary combinations of x, y and z. As expected, the higher the number of events to be detected in a sequence of time intervals, the lower is the number of false positives recorded by the nodes. In the office scenario, this allowed to limit the number of false positives (and reduce to zero for the case with 8 events detected within 12 time periods of 250 milliseconds each). For the extremely noisy environment with a radio running, however, the extended event detection mechanism was able to reduce the number of false positive events, but the remaining number remained significant. This implies that – like for the assignment using light events – a certain control over the conditions during the assignment of symbolic coordinates using acoustic events is required.

Table 4 Extended false positive analysis

	Office environment	Radio running
Single event	336	6068
5x1000ms, detect 3	46	926
6x500ms, detect 2	116	2469
12x250ms, detect 8	0	391

After considering false positive events, the next step was to evaluate the success rate of assigning symbolic coordinates with the help of acoustic events. For this purpose, we distributed four MICAz sensor nodes in a room and performed a set of experiments assigning symbolic coordinates to these nodes. The distance between the sounder and the receiving sensor nodes varied between one and four meters. We experimented with different combinations of x, y, and z. Table 5 gives an overview of the success rate results that were obtained across 15 experiments for each setting.

Table 5 Average success rates for acoustic events

Values for x, y, and z	**Average percentage of nodes assigned correctly**
5x1000ms, detect 3	100%
5x1000ms, detect 4	98.33%
6x500ms, detect 2	100%
6x500ms, detect 4	96.88%
6x500ms, detect 5	3.13%
12x250ms, detect 8	100%
8x250ms, detect 4	93.33%

The results of our experiments illustrate that the coordinate assignment by assisted broadcast using acoustic events works well and that it is possible to achieve success rates in the range of 100% using appropriate combinations of x, y and z. The event detection process works equally well for different lengths x of the event detection time periods. However, the experiments also confirmed our expectation that the tone detection mechanism is not able to detect events in all of the y time periods as shown by the measurement with six event detection time periods of a length of 500 ms: While four events could be detected with an average success rate of 96.88%, the success rate dropped to only 3.13% when requiring a detected event in five out of the six time periods. We mainly attribute this behavior to the specific hardware used in our experiments and its interaction with the implementation of the timers.

We also placed nodes in neighboring rooms to test whether the used acoustic signals are able to propagate across room boundaries and trigger events at these nodes. However, even though the sound was sometimes still hearable in neighboring rooms, none of the nodes detected any events during the coordinate assignment experiments so that no false positives occurred.

For light events it was reasonable to assume that the effect of an event can be detected irrespective of the location of a sensor node, because the ceiling lighting usually covers the complete area of a room. Our acoustic events, however, are generated by the PDA of the administrator and propagate from its location. The strength of the acoustic signal decreases with the distance from the signal source. Consequently, the detection of acoustic events can only work reliably within a certain radius around the administrator. This does not constitute a serious limitation, because the administrator can safely send the signal multiple times from different locations when assigning coordinates in particularly large rooms. Nevertheless, this factor must be considered.

We evaluated the influence of the distance on the coordinate assignment based on acoustic events with a set of controlled experiments measuring the success rate for different distances between sender and receiver of the acoustic events. We placed MICAz sensor nodes on the floor without any obstacles between the sounder node and the receiver nodes and used two events detected in six 500ms timer intervals as the condition for the coordinate assignment (see the first line in Table 5). Up to a distance to five meters, the experiments showed a success rate of 100%. At six meters distance, the success rate fell to 76.7% across 30 experiments. At a distance of seven meters, the success rate was slightly lower at 70.0%. Note that these results strongly depend on the environment and the placement of nodes in real scenarios. This can be illustrated by an additional experiment where we placed sounder and receiver node eight meters apart but at a height of one meter above the floor. Now we achieved a success rate of 96.7% in the coordinate assignment.

The results confirm that the distance between the sender of acoustic events and the sensor nodes receiving the events is indeed a limiting factor that can occur in practice in rooms exceeding a certain size. However, since walking around large rooms and triggering the event multiple times is easy, we do not consider this a serious limitation.

For some of the problems and limitations described in this section, we expect that sounder and sensor hardware specifically developed for the problem at hand could improve the results. However, this would also lead to increased requirements in terms of the sensor node equipment.

6.3 Summary

Summarizing the results of our evaluation, assigning symbolic coordinates by broadcast has shown a good performance when used together with RSSI filtering. If, however, a high accuracy is required, then assigning symbolic coordinates by assisted broadcast is able to provide a reliable solution as has been shown by our experiments.

The assignment by assisted broadcast using light events clearly outperformed the assignment using acoustic events in terms of reliability. Dealing with the light sensors also proved to be much easier than with sounders and microphones whose behavior was difficult to understand in some cases. The main reason why considering acoustic events is still attractive is that it allows to directly trigger the events from the application without explicit action (i.e., manually turning on the light) by the administrator.

7 Conclusions and Future Work

In this paper we have motivated the benefit of using symbolic coordinates in wireless sensor networks and have discussed the advantages and disadvantages of two canonical approaches for the assignment of such coordinates – the individual assignment and the assignment by broadcast. We have then presented a third solution that combines the advantages of both approaches with only minor additional effort for an externally triggered sensor event. The evaluation shows that this approach allows a simple yet reliable assignment of symbolic coordinates to sensor nodes in indoor scenarios. This way, the manual configuration of symbolic coordinates after the deployment of a sensor network is a viable alternative to more sophisticated node localization approaches.

A reliable indoor localization system for sensor networks that is able to determine the position of nodes without requiring user interaction or expensive hardware while only generating a low message and computational overhead is definitely desirable. However, while this is not foreseeable, our solution provides a reliable assignment of symbolic coordinate information to sensor nodes that only generates a low overhead and only requires a reasonable amount of support by the user. Moreover, these pre-configured static sensor nodes could also provide symbolic location information to mobiles devices during the normal system operation as part of a localization solution.

As part of future work we are interested in a combination of sensor networks and building automation systems, which automate the control of different mechanical and electrical systems in buildings. This would make it possible to turn on and turn off the lights in the rooms of a building automatically. This way, we could completely automate the assignment of symbolic coordinates by sending out symbolic coordinate messages and then triggering the light sensors in the respective rooms of the building. We have already done the first steps in this direction by developing a prototype system that integrates sensor networks with a home automation platform [6].

Another aspect that we are actively working on is the grouping of nodes based on sensor data. Instead of relying on sensor stimuli deliberately triggered by the user we are interested in analyzing the sensor data collected by sensor nodes as part of their normal operation. Based on a similarity analysis of the sensor data from different nodes we want to decide which nodes reside together in the same room and group these nodes together.

References

1. C. Becker and F. Dürr. On location models for ubiquitous computing. *Personal Ubiquitous Computing*, 9(1):20–31, 2005.
2. B. Brumitt and S. Shafer. Topological world modeling using semantic spaces. In *Workshop Proc. of Ubicomp: Location Modeling for Ubiquitous Computing*, 2001.
3. P. Corke, R. Peterson, and D. Rus. Networked robots: Flying robot navigation using a sensor net. In *Proceedings of the Eleventh International Symposium of Robotics Research (ISRR)*, 2003.
4. E. Elnahrawy, X. Li, and R. P. Martin. The limits of localization using signal strength: A comparative study. In *Proceedings of The First IEEE International Conference on Sensor and Ad hoc Communications and Networks (SECON 2004)*, 2004.
5. M. Gauger, P. J. Marrón, M. Handte, O. Saukh, D. Minder, A. Lachenmann, and K. Rothermel. Integrating sensor networks in pervasive computing environments using symbolic coordinates. In *Proceedings of the Third International Conference on Communication System Software and Middleware (COMSWARE 2008)*, 2008.
6. M. Gauger, D. Minder, P. J. Marrón, A. Wacker, and A. Lachenmann. Prototyping sensor-actuator networks for home automation. In *Proceedings of the 3rd Workshop on Real-World Wireless Sensor Networks (REALWSN 2008)*, 2008.
7. C. Jiang and P. Steenkiste. A hybrid location model with a computable location identifier for ubiquitous computing. In *UbiComp '02: Proceedings of the 4th international conference on Ubiquitous Computing*, pages 246–263, London, UK, 2002. Springer-Verlag.
8. K. Langendoen and N. Reijers. Distributed localization in wireless sensor networks: a quantitative comparison. *Computer Networks*, 43(4):499–518, 2003.
9. D. Moore, J. Leonard, D. Rus, and S. Teller. Robust distributed network localization with noisy range measurements. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 50–61, New York, NY, USA, 2004. ACM Press.
10. D. Niculescu and B. Nath. Dv based positioning in ad hoc networks. *Journal of Telecommunication Systems*, 22(1-4):267–280, 2003.
11. N. B. Priyantha, A. Chakraborty, and H. Balakrishnan. The cricket location-support system. In *Proceedings of the 6th Int. Conf. on Mobile computing and networking*, 2000.
12. R. Stoleru, T. He, J. A. Stankovic, and D. Luebke. A high-accuracy, low-cost localization system for wireless sensor networks. In *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, 2005.
13. R. Stoleru, P. Vicaire, T. He, and J. A. Stankovic. Stardust: A flexible architecture for passive localization in wireless sensor networks. In *SenSys '06: Proceedings of the 4th international conference on Embedded networked sensor systems*, pages 57–70, New York, NY, USA, 2006. ACM Press.