User Indoor Location System with Passive Infrared Motion Sensors and Space Subdivision

Marios Sioutis and Yasuo Tan

School of Information Science
Japan Advanced Institute of Science and Technology, Ishikawa, Japan
{smarios,ytan}@jaist.ac.jp

Abstract. The use of indoor location information has the potential to enable ground-breaking smart services in a home environment. The objective of this research is to design a relatively inexpensive real time indoor location system which also poses no threat to user privacy. We propose an indoor location system that uses commodity parts such as infrared motion sensors and the idea of space subdivision. By evaluating the state of the sensors involved in a scene, it is further possible to evaluate each unique area and extract those that the user is most likely to be located at. In conclusion, despite a couple of flaws that should be addressed, the proposed system achieves its targets while maintaining an acceptable level of accuracy.

Keywords: indoor location system, infrared motion sensor, space subdivision, smart home.

1 Introduction

Compared to outdoor location systems where GPS is the de facto standard, there is no such standard for user indoor location systems. In the past many solutions have been proposed, based on technologies such as active RFID[5], cameras [4](survey), ultrasound[2], passive infrared sensors (most notably ThiLo[3]) as well as a number of solutions based on the Received Signal Strength Indication of different wireless systems (Wi-Fi[7], Bluetooth[1], ZigBee[6]). It is the authors conviction that such indoor location systems will be an indispensable part of a platform on which other smart services can be developed and assist the user in his everyday activities.

In this paper, a real time indoor location system for users utilizing passive infrared motion sensors that are readily available in the market and space subdivision is proposed. The rationale for such a system becomes quite clear when looking at the advantages of the system.

First, compared to indoor location systems that utilize cameras, the proposed system does not pose a threat to the user's privacy. For many users, the possibility of network enabled cameras being exploited and their streams subverted is enough of a deterrent for installing such systems, further hindering the adoption of smart technologies in the home environment.

N. Streitz and P. Markopoulos (Eds.): DAPI 2014, LNCS 8530, pp. 486–497, 2014.

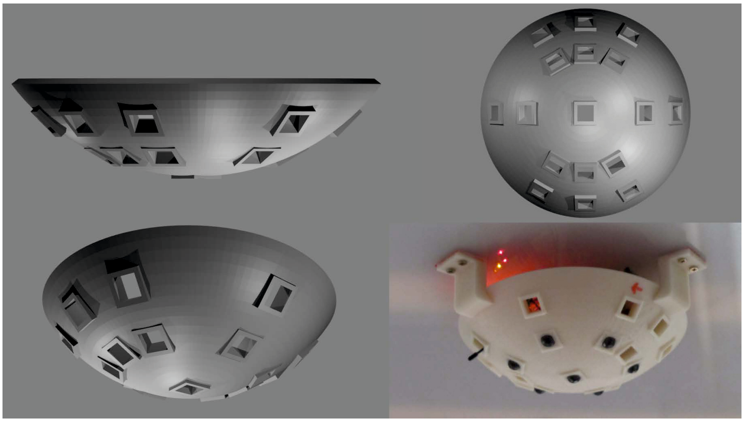

Fig. 1. The sensor pod design and the final 3D printed prototype

Second, in contrast to systems that use active or passive transceivers the user is free from wearing or holding any such transceiver which is in turn one less hurdle towards the deployment of indoor location systems.

Third, due to the design of the sensor pods used in this system, it is unaffected by typical household furniture and the deployment cost is minimal. Furthermore, the cost of each sensor pod prototype can be roughly estimated at 200 USD, a cost that can further be suppressed.

Fourth, the proposed system is extendable and can be easily customized. Extending the coverage area is as simple as adding more sensor pods to the desired rooms. Furthermore, the position as well as the actual shape of the sensor pod can be customized to better fit the peculiarities of any room.

Finally, in contrast to ThiLo[3], no initialization step is necessary.

The proposed system has two significant downsides. First, compared to other solutions based on cameras and active/passive transceivers, the accuracy of the reported user location may be inferior. However, it can be argued that the accuracy of the predictions provided by this system is enough for developing location aware services in the home environment. Furthermore, the current system has no means of deducing the identity of a user.

The remaining sections of this paper discuss the technical aspects of the system and its design, the idea of subdividing space and how it is utilized in the proposed system, sensor behavior, the algorithm used for evaluating areas, the experimental setup and evaluation methodology, points of consideration and finally conclusion and future work.

2 Sensor Pod Design and Technical Considerations

The proposed system uses the AMN3111 passive motion sensor produced by Panasonic as the basic unit of collecting information. This sensor has a wide

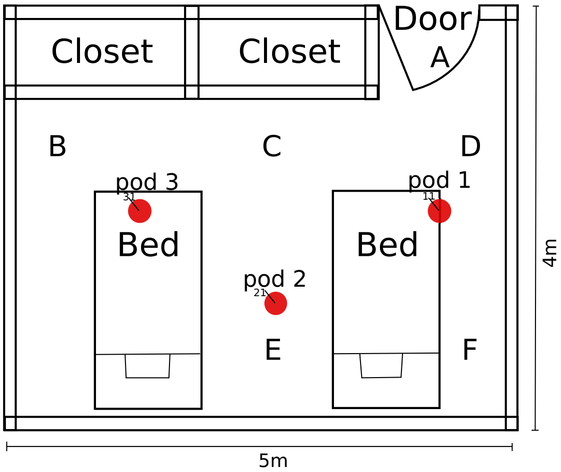

Fig. 2. Room details and sensor pod positioning

sensing angle of 90^o degrees, a maximum detection distance of 5 meters as well as good response times. Our system utilizes these sensors as a cluster of 7 sensors, mounted on a sensor pod. Multiple such sensor pods can then be affixed to the ceiling of a room to provide user location information. The 7 sensors of a sensor pod are positioned as seen in Fig. 1. The first sensor faces downwards towards the floor, whereas the remaining six sensors are set so as to have an angle of 60^o degrees with respect to ceiling plane and are spaced out evenly every 60^o degrees on the half-spherical fixture of the sensor pod.

Three sensor pods where affixed to the ceiling of a room, as seen in Fig. 2. Excluding the first sensor of each sensor pod, the orientation of the second sensor is marked[1] as a short black line, along with the sensor's assigned ID tag. The remaining sensors are evenly spaced out by 60^o degrees.

For controlling the sensors and transmitting data, one Arduino Fio board, an XBEE low power wireless module as well as a custom circuit board for multiplexing the output of the sensors were used. For collecting the data, one Arduino Fio equiped with an XBEE module was connected to a laptop. The raw data gathered for one of the experimental runs can be seen in Fig. 3. A green round mark shows that the given sensor is in a logical "HIGH" state, whereas a red triangle mark signifies a logical "LOW" state. These states are indicative of the presence or the absence of a user.

Finally, visualization of the coverage area and location results was achieved with custom software developed using Java and Open GL ES. The sensor pod design and video processing was done in Blender, a powerful open-source software mainly targeted at 3D modeling.

[1] Also marked with a small red arrow on the prototype unit.

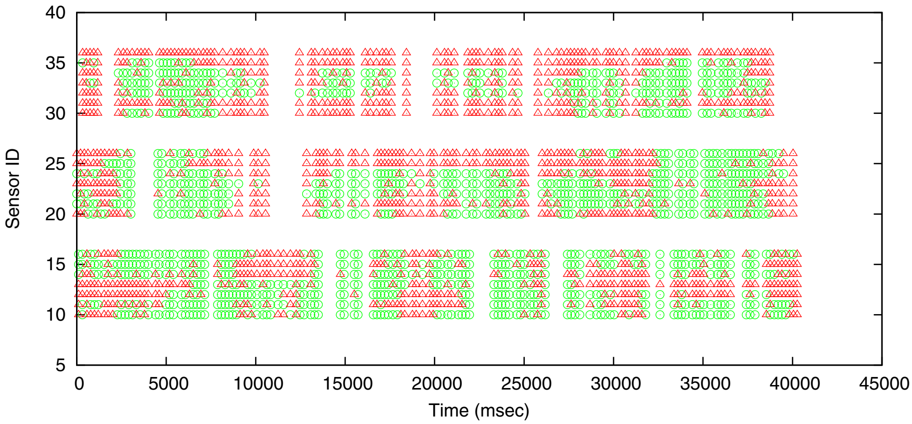

Fig. 3. The timeline of sensor events for the first run

3 Space Subdivision

The design of the sensor pod combined with the wide 90^o degrees coverage angle of the sensors gives rise to some very interesting properties. Sensors 2 to 7 are spaced evenly in a circle pattern with their direction changing by 60^o degrees with respect to the elevation axis, thus creating an overlapping area of 30^o degrees between two consecutive sensors. Furthermore, the existence of the first sensor which is pointing downwards further subdivides space in areas that are "close" to the actual sensor pod and those that are further away. From the above, it is obvious that even with a single sensor pod useful deductions regarding the user's position in relation to the sensor pod can be made. The coverage area of a single sensor pod can be seen in Fig. 4 (left).

Further increase of the number of sensor pods in a scene leads to a rapid growth of the number of unique areas in which the scene space can be subdivided. During experimentation, the number of unique areas that were reported for the experiment scene used in the evaluation ranged from 700 to 850 unique areas. The number of unique areas varies depending on the coverage angle of the sensors, the location of the sensor pods as well as the furniture in the scene.

The octree used to model the scene and the coverage area of the sensors is 8 levels deep, with the nodes of the deepest level having a cube size of only 4 cm^3. By keeping track of which sensors cover which nodes, it is possible to extract the overlapping regions among sensors, thus creating unique detection areas for a given combination of active sensors.

One more topic worth mentioning is the number of sensors each unique area is covered by. Most areas are covered by anywhere between five and fourteen unique sensors, as shown in Fig. 4 (right). In this figure, areas that where covered by less than five sensors have been pruned from the scene, as such areas were marginal and very close to the ceiling, thus not corresponding to a realistic position for the user. The progression and change of the colors signifies the increase in the

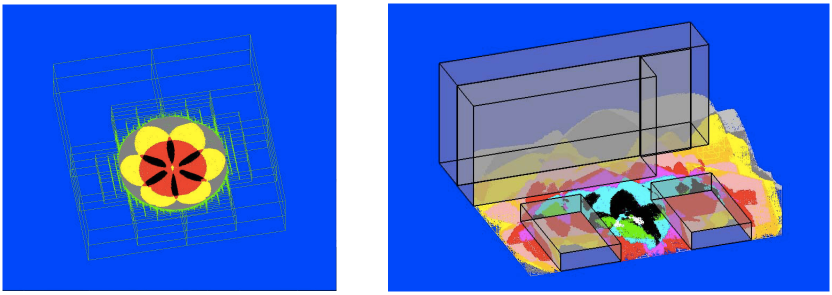

Fig. 4. Left: coverage area of a single sensor pod. Right: combined coverage area of a room with three sensor pods.

numbers of sensors an area is covered by. Starting from the edges of the room, light grey and gray areas are covered by five and six areas respectively. Moving progressively to the center of the room, areas colored black, green and white are covered by thirteen, fourteen and more than fourteen sensors respectively.

4 Sensor Behavior

4.1 Behavior Characteristics of an Ideal Digital Sensor

Before we look further into the specifics of how the infrared motion sensors that were deployed in this system work, the case for an ideal digital sensor must be made. Such an ideal sensor is able to report whether a user can be detected in its coverage area, even if the user is partially in this area. The absence of the user should also be reported as soon as the user leaves the area. Furthermore, such an ideal sensor must always report that a user is detected, even if the user stops moving entirely. The presence or absence of a user are communicated as a digital output of "HIGH" and "LOW" respectively.

As part of testing, sensor data from such an ideal sensor were generated by representing the user as a rectangular cuboid box that is 1.75 meters tall and 40 centimeters wide. This virtual user moved inside the modeled scene and the sensor data that was generated was fed back to the visualization system. As expected, the location prediction transitioned smoothly from one location to the next with great accuracy, verifying the core idea of the proposed location system.

Despite the arguably favorable results that such a sensor would produce, to our knowledge such a sensor has yet to be invented. The actual sensors used in this system have somewhat different behavior than the ideal sensor, which leads to certain complications. The characteristics of the sensors deployed will be discussed in the next section.

4.2 Behavior Characteristics of the AMN3111 Sensor

In stark comparison to the ideal sensor, the behavior of the AMN3111 sensor differs greatly in two aspects. First, the time this sensor needs to report the absence of a user exiting its coverage area varies from a couple hundred milliseconds to as many as a couple of seconds, depending on the distance from the sensor and the relative angle the user is exiting the region from.

The second aspect in which the AMN3111 sensor differs from the ideal digital sensor is the behavior regarding a user that remains relatively motionless inside the sensor coverage area. In such cases, the AMN3111 will stop reporting the presence of the user inside the area after a couple of seconds (or even longer, again depending on the relative position of the user to the sensor).

Both of these behavioral differences produce undesired side effects and interfere with the location prediction algorithm and should be accounted for. Their exact effects as well as the countermeasures that can be taken to avoid them are described in Sect. 7.

5 Algorithm Explanation

The proposed algorithm can be split into two parts, a sensor-independent area evaluation strategy and a sensor-dependent sensor state evaluation part. These are described in the next two sections.

5.1 Area Evaluation Strategy

The sensor-independent part of the proposed algorithm deals with the evaluation of each individual area. This part depends on the deduction of each sensor state and its assignment of a numerical value by the sensor-dependent part of this algorithm.

The sensor-independent part uses a set of key-pair values. Each entry in this set corresponds to a sensor (key) and the actual assigned numerical value for its given state (value). The sensor independent part then proceeds to compare this set of key-value pairs against each unique area in the scene.

At the start of the algorithm, the score for an area is zero. For each key-value pair, if the sensor is covering that unique area then the score for this area is increased by that sensor's current value, else that amount is subtracted. The top 2% of the areas with the highest scores is then visualized.

5.2 Sensor State Evaluation

The sensor-dependent part of this algorithm deals with the detection of the state the sensors in the current scene are in and the assignment of numerical values to these states.

The current implementation defines seven states, *RISING*, *UPTRIGGER*, *UP*, *FALLING*, *DOWNTRIGGER*, *DOWN* and *FLUX*. The first three states are associated with the presence of a user in its coverage area. In the *RISING* state, the sensor transitioned recently from a "LOW" state to a "HIGH" state in the past 300 msecs and has a numerical value of 30. In the *UPTRIGGER* state, the sensor transitioned to a "HIGH" state more than 300 msecs ago, thus it can be deduced with fair certainty that a user has indeed entered the area recently. Its numerical value is 100. The *UP* state shows that the sensor has been in the "HIGH" state for more than a second continuously. It has an associated value of 50.

The next three states are the logical equivalents of the first three states with the only difference being that the sensor transitions from a logical "HIGH" state to a logical "LOW" state, thus detecting the absence of a user. These states are assigned negative values equal to those of their equivalent positive states.

The *FLUX* state shows that in regards to this sensor, three (or more) discrete and logically alternating events have occurred in the past second, thus the sensor is in a state of flux, unable to deduce whether a user is in its coverage area or not. This state is associated with a numeric value of five, a slightly positive value, because the user is probably located at the edges of this sensor's detection area. It is not uncommon for a sensor to transition to a flux state first before it is possible to reliably determine its state.

This part of the algorithm also assigns a certain percentage bonus to the sensor score, depending on the time elapsed since the sensor entered its given state. The sensors are chronologically sorted and the sensors that most recently transitioned to logically "HIGH" states may get up to a hundred percent bonus to their base score, with the bonus assigned dropping linearly. Similar logic applies to the sensor states associated with user absence.

This part of the algorithm is sensor and circuit dependent in the sense that, other passive infrared motion sensors may exhibit different behaviors that can possibly be expressed better with a different time scale or a different set of states and transitions. Although our conviction is that the proposed states reflect the behavior of the AMN3111 sensor well, this may very well not be the case for other sensors.

6 Experiment Setup, Evaluation and Results

6.1 Experimental Setup

The experiments described later in this section where conducted in a single room as seen in Fig. 2. The room is 5 meters long and approximately 4 meters wide. The ceiling is 2.4 meters tall. The modeling of the room was based on exact floor plans and the space taken up by the wardrobe and beds is properly taken into consideration and removed from the scene's octree.

The positions of the sensor pods was chosen so that they would create an almost equilateral triangle, with one of the sensor pods slightly displaced to the right. The reasoning behind this formation is that it would create the most evenly distributed unique areas. Furthermore, the slight displacement of the rightmost sensor pod allows for better coverage of the door area.

In total, four recorded runs were performed, each one filmed by one camera view point. The paths taken for each of these runs can be seen in Table 1. A parenthesis around a letter means that the user passed through that given point

Table 1. Run paths

Run Name	Run Points
First run	E, (C), (D), A, (D), F, D , C, E
Second run	B, (C), D, (C), B
Third run	B, (C), D, A, D, (C), B
Fourth run	F, (D) , A, (D), F

without stopping, otherwise the user briefly stopped at that point for a brief amount of time (2 to 4 seconds).

In the first run, the user takes a more complicated path with more temporary stops compared to the other runs. Runs two and four were selected to test the accuracy of the system for straight paths; in run two the user walks along the length of the room, whereas in run four the user walks along the sort side of the room towards the door. Run three is a combination of run two and four, where the user walks along the length of the room, but makes a right turn towards the door. In all of the runs the user returns to its original position.

6.2 Evaluation

To evaluate the accuracy of the system, a qualitative comparison between the output of the visualization software and the actual footage is performed. For this purpose, virtual camera points were set up in the visualization software that matched the orientation, position and field of view of the original camera used to take the actual footage. Then, using the blue color as a channel key, transparency is added and the visualization output is superimposed on the original video.

For comparison purposes, the resulting video is also rendered as still pictures, which were used for evaluation. For the first run, ten frames per second were used. Eventually, this was deemed to not significantly increase the accuracy of the results, and for the second, third and fourth runs three frames per second of running video were considered adequate.

Each frame was assessed qualitatively in a five-rank scale, from excellent (five points) to bad (one point). The definition for each grade is as follows:

- excellent: the reported location matches and it is tightly focused around the user,
- good: the reported location matches, with a small percentage of it pointing out towards another direction
- fair: the reported location matches but it is relatively broad and/or visualized as floating, or the user is slightly outside the reported location (less than 50 cm) but the area still remains focused
- poor: the user is clearly outside of the reported location but less than a meter away
- bad: the user is clearly more than a meter away from the reported location, or the reported location is too wide and/or fragmented to be meaningful.

An example for each of the above categories can be seen in Fig. 5

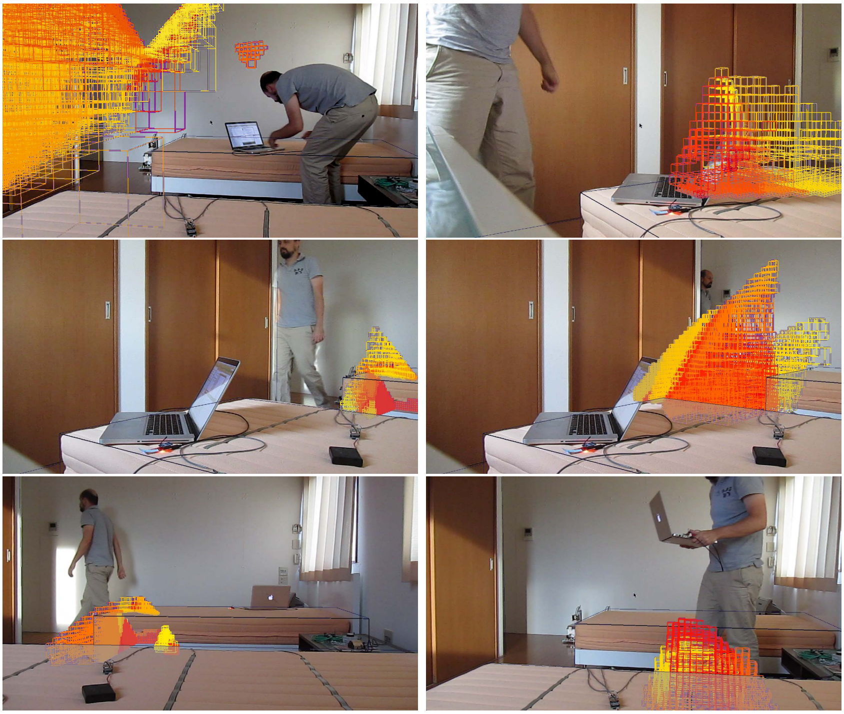

Fig. 5. Examples of evaluation frames.
Top left: a bad frame. Top right: a poor frame.
Middle left: a fair frame (near miss). Middle right: a fair frame (large area).
Bottom left: a good frame. Bottom right: an excellent frame.

Table 2. Experimental results

Run	Total Frames	Avg. Score	Bad	Poor	Fair	Good	Excellent
first	497	3.62	22	90	107	112	166
second	57	3.07	3	17	17	13	7
third	88	3.11	13	11	31	19	14
fourth	74	3.20	6	11	32	12	13

6.3 Run Results

The results of the four runs can be seen in Table 2.

The results are broken down to the total number of frames evaluated, an average score as well as the number of frames for each of the five categories. What is not readily apparent from Table 2 is the fact that most of the 'bad' states occur close to the start and end segments of each run. This can be seen clearly in Fig. 6, which represents the timeline for the state detection of the first run. 14 out of 22 bad states occur near the beginning of the run and another 4

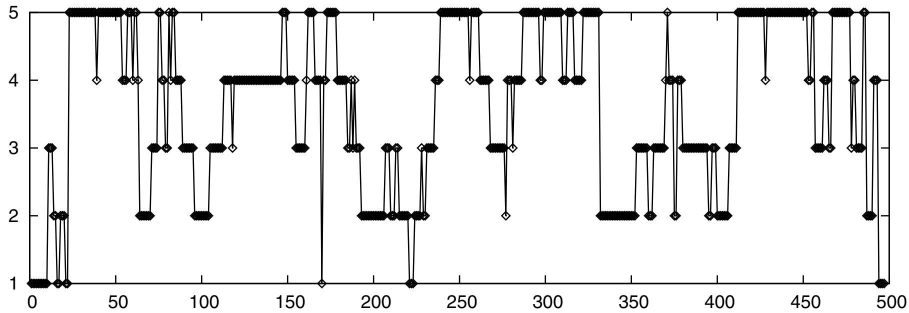

Fig. 6. State detection timeline for the first run

bad states occurring at the end of the run. Thus, only in 4 out of the 497 frames the actual reported user location was 'bad'.

On average, the system achieved a score more than 3 in all performed runs. This indicates that the average quality of the reported location should fit somewhere between the fair and good scales explained earlier. As a result, it is not a stretch to claim that the inaccuracy of the positions reported by the system is less than 50 centimeters.

7 Points of Consideration

There are two types of undesired side effects that manifest in this system, and the cause for both of them can be linked to the behavior the AMN3111 sensors exhibit. These side effects are explained in the next two subsections.

7.1 Strong "Gravitational" Field Effect

The first side effect is the presence of what can be casually described as a strong "gravitational" field directly below the installation point of each sensor pod. When the user walks away from such a gravitational field, the location projection the system reports has the tendency to lag behind the actual user location, seemingly unable to escape the "gravitational" field. This phenomenon is very pronounced in the second run as well as some parts of the forth run.

The reason for this phenomenon is the time delay observed between the user leaving a specific area and the actual sensor reporting this absence, as mentioned in sect. 4.2. When the user leaves such an area, some sensors linger longer to a "HIGH" state thus creating false positives. As a consequence, during the evaluation step of each area, the areas in the "gravitational" field that also happen to be covered by a larger number of sensors tend to be more favorably evaluated, earning higher scores. The bonus scoring strategy for the sensor does help alleviate this problem to a small extent, by assigning very small bonuses to areas that have been constantly on for a long time and awarding higher scores to sensors that transitioned to a logical "HIGH" state more recently. However, it proved difficult to completely alleviate this side effect.

Another possible solution to this problem is the implementation of a strategy to decide the most optimal shape and position for the sensor pods, so as to reduce the existence of such heavily covered areas, thus avoiding such "gravitational" fields.

7.2 Relation between Bad States and Loss of Information

Bad states that were detected in the four experimental runs are all associated to loss of information. The nature of this loss becomes clear when we consider the behavior of a sensor while the user remains motionless inside its coverage area: the sensor will eventually transition to a logical "LOW" state. This has the exact opposite effect from the previously described strong "gravitational" field: areas that are covered by multiple sensors now are heavily penalized as more and more sensors turn off, and the system favors marginal areas at the sides of the room. Such areas are covered by only a few sensors and are less heavily penalized when the number of sensors that go off increase.

As a counter measure, in the four runs presented earlier, a simple method was used: as soon as a sensor pod would report less than two active sensors, the state of the sensors for that sensor pod was frozen in time, representing a "last known good state". This technique served well to alleviate this problem when the user stands still during a run. Ultimately though, when a run comes at its end, all sensors transition to a logical "LOW" state. At this point it is not uncommon that less than six sensors are frozen in a logical "HIGH" state. In conjunction with the fact that most unique areas are covered by significantly more than six sensors, this loss of information prevents the system from making an accurate prediction when it comes to static users.

An algorithm that solves this end state problem can be implemented with relative ease. Such an algorithm would freeze the user location prediction as soon as the number of active sensors in a scene runs below a certain cut-off point. This cut-off point should take into consideration the distribution of the number of sensors the unique areas are covered by and adjust accordingly.

8 Conclusion and Future Work

In this paper, an indoor location system that utilizes the AMN3111 passive infrared motion sensors and subdivision of space was proposed. A qualitative evaluation of the system for four experimental runs showed that the system is fairly accurate and on average it reports a user location that reflects reality with possible small deviations.

The system compares favorably when compared against other location systems that do not use motion sensors in terms of user privacy, ease of use (no passive/active tags or badges are needed), extensibility, customization, deployment and cost. In comparison to ThiLo[3], there is no need for an initialization step. However, other solutions are capable of more accurate predictions and user identification.

As shortcomings of the system, two side effects were reported: strong "gravitational" fields and the relation between end states and loss of information. Both problems are closely related to the behavior of the sensors used in this system.

As future work, further research needs to be conducted in order to solve the two previously mentioned problems. Furthermore, alternative area evaluation methods such as particle filters should be evaluated and contrasted with the current implementation. Furthermore, a great challenge for this system would be to report the location of more than one user in a single room simultaneously. Lastly, an auxiliary user identification mechanism that would be used in conjunction with this system should be pursued.

References

1. Feldmann, S., Kyamakya, K., Zapater, A., Lue, Z.: An indoor bluetooth-based positioning system: Concept, implementation and experimental evaluation. In: Zhuang, W., Yeh, C.H., Droegehorn, O., Toh, C.T., Arabnia, H.R. (eds.) International Conference on Wireless Networks, pp. 109–113. CSREA Press (2003)
2. Harter, A., Hopper, A., Steggles, P., Ward, A., Webster, P.: The anatomy of a context-aware application. In: Proceedings of the 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking, MobiCom 1999, pp. 59–68. ACM, New York (1999), http://doi.acm.org/10.1145/313451.313476
3. Hauschildt, D., Kirchhof, N.: Advances in thermal infrared localization: Challenges and solutions. In: 2010 International Conference on Indoor Positioning and Indoor Navigation (IPIN), pp. 1–8 (September 2010)
4. Mautz, R., Tilch, S.: Survey of optical indoor positioning systems. In: 2011 International Conference on Indoor Positioning and Indoor Navigation (IPIN), pp. 1–7 (September 2011)
5. Ni, L., Liu, Y., Lau, Y.C., Patil, A.: Landmarc: Indoor location sensing using active rfid. In: Proceedings of the First IEEE International Conference on Pervasive Computing and Communications, PerCom 2003, pp. 407–415 (March 2003)
6. Sugano, M., Kawazoe, T., Ohta, Y., Murata, M.: Indoor localization system using rssi measurement of wireless sensor network based on zigbee standard. In: Fapojuwo, A.O., Kaminska, B. (eds.) Wireless and Optical Communications. IASTED/ACTA Press (2006)
7. Xiao, J., Wu, K., Yi, Y., Ni, L.: Fifs: Fine-grained indoor fingerprinting system. In: 2012 21st International Conference on Computer Communications and Networks (ICCCN), pp. 1–7 (July 2012)