PREDICTIVE SIMULATION & EVACUATION MONITORING IN WIRELESS SENSOR NETWORKS

Seán Óg Murphy[12], Kenneth N. Brown[2], Cormac Sreenan[1]

1. Mobile and Internet Systems Laboratory, 2. Cork Constraint Computation Centre,
Department of Computer Science, University College Cork, Ireland
jmm3@student.cs.ucc.ie, k.brown@cs.ucc.ie, cjs@cs.ucc.ie

Abstract. In this paper we present an emergency monitoring strategy (Predictive Simulation Reporting) using Wireless Sensor Networks that combines sensor monitoring and the prediction of future movement of occupants using a multi-agent crowd evacuation simulation operating on a resource-rich Base Station node. Sensor nodes are informed ahead of time about multiple possible evacuation scenarios that could occur in the near future. As the evacuation proceeds, sensor nodes compare the candidate scenarios with their actual readings to determine at an early stage where the occupants are going. Our preliminary experiments indicate that this strategy allows the nodes and base station to identify a matching scenario quickly and efficiently.

Keywords: wireless sensor networks, evacuation monitoring, multi-agent simulation

1 Introduction

The safety and speed of emergency evacuation of buildings can be improved with the exploitation of networked “smart building” surveillance and evacuation control components [1] and centralized evacuation planning algorithms[2]. Monitoring of smart buildings is often performed using wireless sensor nodes which are battery-powered and have limited computation capability. In the emergency evacuation scenario, wireless sensor networks (WSNs) monitor the occupancy of the building and location of hazards (fire, smoke, chemical spill etc.) and this information is combined with knowledge of the building layout to generate evacuation plans. The plans can be computed at a centralized Base Station using an optimal evacuation plan algorithm based on initial occupancy, knowledge of the building structure and knowledge of the behaviour of the hazard. The evacuation plan can be enacted in the building in the form of a set of instrument actuation instructions, such as direction signs configured to direct occupants along the optimal route.

Timely response to events in a building during emergency is of key importance to ensure sensible evacuation planning; if egress of the building proceeds in an unexpected manner, the system needs to respond quickly and generate an updated evacuation plan to ensure the safe evacuation of the occupants. Network communication combined with computation of updated plans may lead to plans which are already out of date by the time they are received by nodes in the network.

Our goal in this research is to reduce the communication burden on sensor nodes and improve response time to changes in evacuation progress by using a resource-rich Base Station (BS) node to predict future movement of building occupants. We use a Multi-Agent Crowd Evacuation simulation at the BS node to simulate several possible future progressions of the evacuation. The BS simulates how each of these candidate scenarios would be perceived by the sensors and computes the optimal evacuation strategy in each scenario, and then sends this information to the wireless sensor nodes. As the evacuation proceeds, the sensor nodes use this information to identify the simulated scenario that best matches their actual readings. If the sensors determine that one of the candidate scenarios matches closely to their readings, they can then inform the nearby direction signs as to the proper evacuation plan to enact without requiring a time-consuming multi-hop report to the Base Station.

Our contribution is a novel combination of a multi-agent-based simulation tool with a monitoring methodology that operates in real-time to improve local response time to changes in progress of the building evacuation. In this paper we will describe some related work in predictive WSN monitoring, followed by a description of the EvacSim Evacuation simulation tool. We then present the methodology behind candidate event simulation at the Base Station and the sensor trace matching procedure at the sensor nodes. This is followed by the description and results of some preliminary experiments performed to evaluate the operation of this evacuation monitoring strategy. Finally we conclude and outline future research directions.

This work was funded by HEA PRTLI IV as part of the Nembes project.

2 Related Work

Goel and Imelski's work on Predictive Monitoring (PREMON) [3] demonstrated the energy savings possible when future sensor readings are predicted centrally, and they show benefits of exploiting correlation between readings on different sensors while minimizing work performed by the sensors. PREMON is motivated by the desire to reduce the energy consumption of wireless nodes and utilises an MPEG-style encoding scheme which uses a global view of sensor reading history to predict the next global set of sensor data. The predicted values are sent to sensors and the sensors need only report back to the Base Station if the predicted values mismatch with the real readings. This approach predicts future readings based on previous sensor readings (as “frames” in an MPEG-style “movie”). However, PREMON generates only the most likely future prediction. In scenarios with multiple tracking targets, such as a building evacuation, there are many possible future events, and there is a low possibility of being able to predict the true future event. Predicting the wrong event, and thus actuating the wrong guidance, could have a serious negative effect.

The Dual Prediction-based Reporting (DPR) mechanism [4] operates using duplicate predictive models running on the base station and on individual sensor nodes as a basis for single-target tracking. Each node uses the model to predict future sensor readings based on historical data. As the base station possesses duplicates of these predictive models, the sensors need only inform the BS of their readings in the event that the predictive model failed to match with the observed readings. The predictive model is based on linear-projection of previous target positions and sensors share their historical readings with neighbours to facilitate the tracking of the target as they move from one sensor field into another. DPR demonstrates significant energy savings over constant monitoring and PREMON but suffers from the poor scalability of its predictive model. Tracking of large numbers of individuals requires a large set of predictive models at each sensor node to predict future positions of each individual. DPR also requires that the sensors are capable of maintaining knowledge of the identity of individuals in the building: the predictive model predicts future movement of an individual based on previous known positions of that particular individual.

The Predictive Simulation Reporting mechanism used in our work uses a multi-agent simulator (EvacSim) similar to the “Distributed Building Evacuation Simulator” [5]. DBES is an emergency evacuation simulation tool intended to evaluate evacuation planning and resource allocation strategies in "smart" buildings (buildings that contain embedded sensor networks and signpost instrumentation. Such a simulator is required as it is impractical to evaluate these strategies using real-world tests such as fire drills. DBES is designed to run offline in either a centralized or distributed manner and uses a graph-based mobility model for agents and hazards. Other research in multi-agent crowd simulation for emergencies include work by Korhonen et al [6], Tavares and Galea [7] and Pelechano and Badler. [8]

3 Simulation for Event Prediction

3.1 Assumptions

In our work we assume that a structural model of the building is available and that the field of vision and duty cycle of sensor nodes is known by the Base Station. We assume, before the emergency is reported, that the Base Station has built up a perspective of the building occupation density in each room, based on reports from the sensor nodes. We assume that an emergency alert is generated by the network (for instance, a smoke alarm triggers as a result of fire), and that this alert informs the Base Station as to the location and likely spread of the hazard. We assume that the WSN includes signpost actuation nodes which can be instructed via wireless communication to direct occupants along particular routes. We assume that the nodes have timekeeping capabilities and that the Base Station and wireless nodes are loosely synchronised. We do not assume that there is full sensor coverage of the building.

Progress of the evacuation may evolve in a number of ways, for instance occupants might all choose to head towards a particular exit, or a group of individuals might ignore the signpost guidance and take an unexpected route to safety. Occupant movement is limited by physical constraints such as inability to pass through obstacles, and is motivated by typical human behaviours such as forming groups that travel together, and avoiding danger.

In our approach (Predictive Simulation Reporting) the BS generates multiple possible futures ("scenarios"), covering short periods of time (e.g. 60 seconds) and produces a simulated set of sensor readings for each scenario. Each of these scenarios represents different combinations of routes occupants could take to safety, and variations in occupant response to the signpost actuation. The simulated readings are sent to the sensor nodes in the WSN ahead of time and the surveillance sensors are required only to compare their readings with the simulated readings to determine the best matching candidate scenario. The optimal evacuation plan for each scenario is sent in the form of actuation instructions to actuation nodes in the network (e.g. via multi-hop wireless communication)

Our aim is for the sensor nodes to detect if the evacuation is progressing according to one of these candidate scenarios and inform the actuation nodes as to which evacuation plan they should implement. We also wish to reduce the communication burden for reporting evacuation progress to the Base Station by requiring the sensor nodes to simply report the degree to which each candidate scenario matches what they actually observe, rather than detailed periodic sensor readings. We believe that this approach will improve the response time in the network to changes in evacuation progress without significantly increasing computation or communication performed by the wireless sensor nodes.

The candidate scenarios are simulated using a multi-agent evacuation simulation (EvacSim) which simulates sensors and building occupants, generates the scenarios and produces these simulated sensor readings. The scenarios are generated based on initial occupancy data and different evacuation routes that groups of occupants in the building might take. For each such candidate scenario, a simulation cycle is performed using the EvacSim simulator. Each simulation cycle produces a set of simulated sensor readings (based on observations of the simulated sensors) which is associated with that candidate scenario.

3.2 Candidate Event Confidence Calculation

Each sensor in the real network receives their set of predicted sensor traces and ID of the associated candidate scenario for each trace. Predicted sensor traces are strings of binary values representing a time series of predicted readings. During the real evacuation, each sensor periodically compares their observations with the expected values for each candidate event and determines their confidence in any particular event being a match for the real event.

In the event of a good match, the sensor can report the ID of the best match to the Base Station. In the case of no good match, the sensor reports their historical readings to the Base Station which updates its view of the building based on the unexpected progression of the evacuation.

In this paper we assume that the occupant monitoring sensors are a simple binary motion detector with limited range. Such sensors detect the presence or absence of occupants within their range as a simple Boolean True/False value. A Sensor trace for these sensors is in the form of a time series of True/False values. These sensors are unable to determine the total number of occupants they can sense and cannot identify specific individuals; they merely detect presence of occupants within their field of view. This model was chosen as more sophisticated sensor types can be interpreted with this binary model (e.g. a complex video camera can report binary True/False if motion is detected).

The sensors compare their real readings with the predicted traces by determining the number of positive readings that appear in the predicted trace and the real readings. This is performed as a simple AND operation over the two strings of Boolean reading values and the total number of matches represents the matching value of the simulated event. Each candidate scenario is given a value based on the number of matches and the event with the greatest number of matches is chosen as the best matching scenario.

Confidence assessments are performed periodically on each sensor, taking all of the sensor readings since the last assessment and determining the total number of positive matches for each candidate scenario. Each such assessment can be reported to the Base Station as a sensor report, and if an individual sensor is particularly confident in a scenario (where that scenario has the highest confidence value and this value is greater than the next highest by a given threshold), they can instruct nearby signpost nodes to implement the associated evacuation plan.

As confidence is based on the proportion of matching non-zero values, this approach is relatively robust to small changes in the start and finish of positive readings (e.g. occupants arriving late at a sensor). While the observations in a reading sequence might not perfectly synch with the scenario trace, there is still some overlap that generates positive reading matches and increases confidence in the scenario.

4 EvacSim Evacuation Simulator

To generate the scenarios, we have constructed EvacSim, a Java-based building evacuation simulation tool, illustrated in Figure 1. This tool models a building in 2.5 dimensions (2D floors stacked on top of each-other) and represents building occupants as individual agents with dynamic movement behaviour (flocking, group formation, obstacle avoidance) and path planning capability. Occupant Agents move through in the building according to their plan, can respond to visible signposting or hazards and form groups that move and plan travel together. Occupant Agents travel in 2-dimensional space, directing their travel through manipulation of a 2D vector which they use to steer towards their goal while avoiding obstacles.

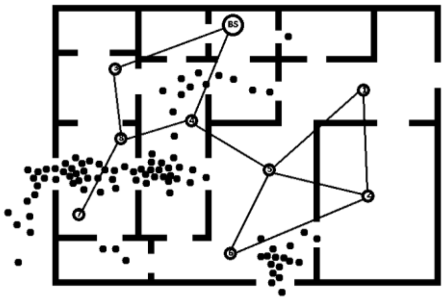

Figure 1: EvacSim Simulator

EvacSim also models individual sensor and actuation nodes as simulated WSN elements. Simulated sensors model the field of sensor coverage, sensing accuracy and duty cycle of equivalent sensors in the building. Actuation nodes are modelled as doorway signposts which can be activated to direct occupants towards or away from entranceways. The sensors and actuation nodes are obscured by building geometry: for instance a sensor cannot detect occupants through walls nor can an occupant observe signposting hidden from view.

Simulated sensors observe the occupant agents in each scenario and produce a trace of simulated readings associated with that scenario. These traces can be collected after a simulation cycle and pushed to the nodes in the real network.

5 Experiments

Two sets of experiments were performed to investigate the feasibility of the Predictive Simulation Reporting scheme. The first experiment investigates the time taken for individual sensors to determine a strong positive match to a scenario, and the number of correct and incorrect matches given by the sensors; the goal of this experiment is to determine the degree to which individual sensors can be relied up to identify the matching scenario based on their own observations. In this experiment, the sensors identify a scenario match if it has at least 1.5 times more confidence in it than the next best scenario. The second experiment investigates the quality of matching when the Base Station receives the reports from sensor nodes and combines them together.

The experiments were performed in a simulated building (illustrated in Figure 2) with 7 rooms and an open common hallway space. A group of 20 individuals begins the experiment in the top-left room. 12 simulated motion-detecting sensors were placed in the building in a uniform grid. This deployment consisted of a mix of indoor and outdoor sensors in order to monitor the movement within the building and also observe occupants successfully exiting the building.

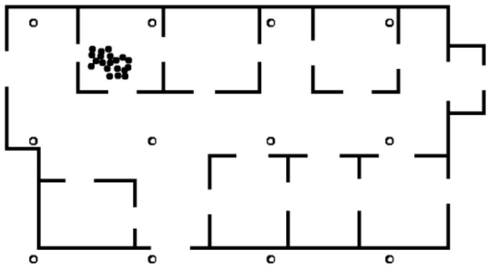

Figure 2: Experiment Building Layout showing initial occupancy and sensor locations

In these experiments, 9 candidate scenarios were generated by directing the crowd to travel to one of 9 destination coordinates (e.g. bottom-right room, top-right exit, top-left exit, top-right room etc.). A 10th scenario was also generated, representing the "real" event, with the simulated sensor readings for this scenario acting as the real readings for the purposes of these experiments. Occupants in the 10th scenario were directed to travel to the bottom-right exterior of the building, which is approximately the same destination as for Candidate Scenario #1.

The sensing duty cycle of the sensors was set to 6 scans per second, with a trace comparison performed once per second. The experiments covered a 10 second period. These experiments were repeated varying the Range and Accuracy of the motion-detector sensors. Range dictates how close an agent needs to be to the sensor in order of the sensor to detect an occupant (where 100 units is approximate to 3 metres) . Accuracy dictates the chance for the sensor to fail to detect an occupant within its field of vision and range (0.0-1.0 where 1.0 = 100% chance).

5.1 Results

The experiments were performed using sensor Range parameters of 80 and 150 units, and sensor Accuracy values of 0.7 and 1.0. For the first experiment, each time a trace comparison is performed the individual sensors record if they have made a scenario match or not, and they note the ID of the matching scenario. If the ID is not that of Candidate Scenario #1, we consider the sensor to be mistaken (they have matched to an incorrect scenario).

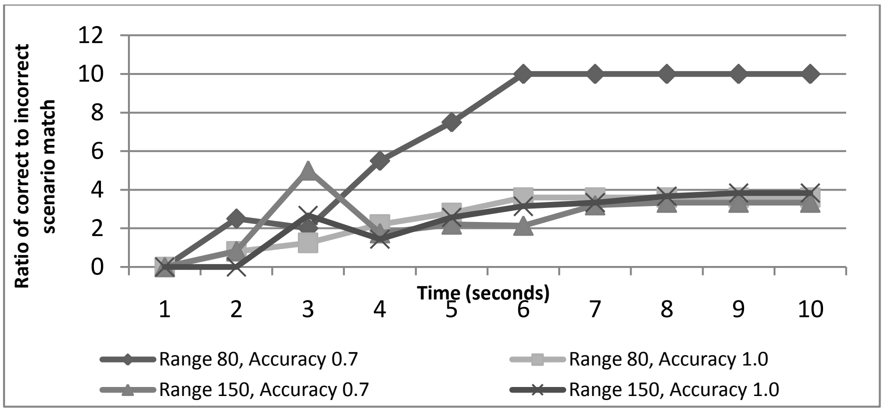

Figure 3: Ratio of correct to incorrect scenario matches at individual sensors (average over 10 iterations)

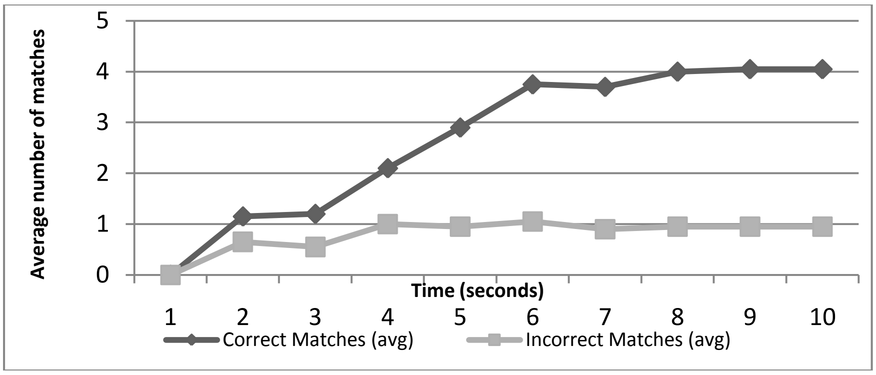

Figure 4: Average number of correct and incorrect scenario matches by individual sensors (average over 10 iterations)

Figure 3 illustrates the ratio of correct to incorrect scenario matches (averaged over 10 iterations of the experiment) made by individual sensors for each combination of Range and Accuracy values in the experiments; for instance with Range of 80 and accuracy of 0.7, after 4 seconds we observe an average of approximately 6 correct matches for every 1 incorrect match.

Figure 4 illustrates the number of correct and incorrect matches over time, averaged over 10 iterations of the experiment: it was found that the average number of correct matches ranged from 0-4 with the number of incorrect matches ranging from 0-1. We observe that the more sensor comparisons have been performed, the greater the number of correct scenario matches. We also observed that a relatively small number of sensors (5-6 out of the 12 sensors) tend to identify no particular scenario, this is a result of the other sensors having a limited view of the "real" scenario and hence lacking enough observations to consider any particular scenario to be a good match.

The second experiment was performed to determine how well the Base Station could identify a matching scenario when it receives all of the Trace Comparison confidence results as periodic report messages from each sensor node (nodes report the values for all the scenarios, not just their best match). By summing all of these scenario confidence values and dividing by the number of Candidate Scenarios, the Base Station can determine if any particular scenario is a good match.

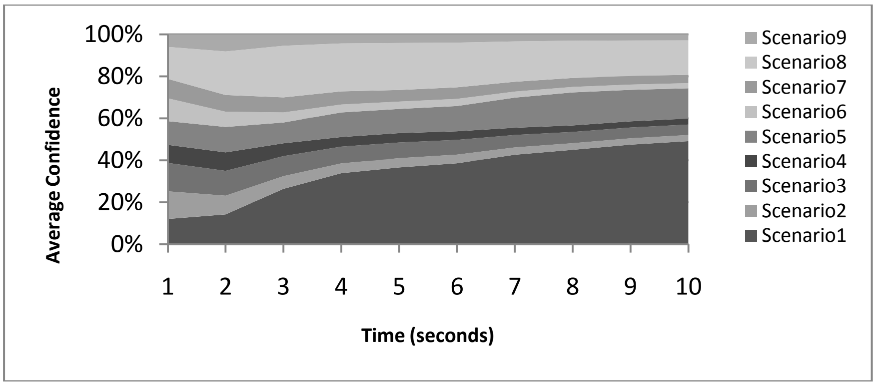

Figure 5: Average % confidence in each scenario from Base Station perspective

Figure 5 illustrates the averaged results from each combination of Range and Accuracy values for this experiment. This graph plots the proportion of non-zero matches associated with each scenario as the evacuation progresses. Initially the reports from the sensors paint an ambiguous picture as there is little to distinguish one scenario from another. However, as the crowd begins to move through the building, the average confidence in Scenarios 1 and 8 begins to rise. At approximately the 6-second mark, confidence in Scenario8 begins to fall as the occupants have opted to head to the bottom-left rather than the bottom-right of the building in this candidate scenario. From this point on, the Base Station observes that Scenario1 has at least twice as many non-zero matches with the real data than the next best scenario

6 Conclusions and Future Work

We designed and implemented an emergency monitoring strategy based on prediction using a multi-agent simulation. Experimental results showed that this strategy could be used to achieve a quick response to changes in the progress of an evacuation without requiring terminal nodes to wait for a response from a centralized Base Station. These preliminary results used simplified sensor models and a simple scenario confidence calculation mechanism to identify good scenario matches. We believe that a quick characterisation of the evacuation progress allows nodes to locally make a decision as to the most applicable (pre-computed) evacuation strategy. In future work we will expand on these experiments, implementing more realistic sensor types and investigating more sophisticated scenario matching methodologies at the sensor and base-station.

When individual sensors were required to match their readings with a candidate scenario, we discovered that a relatively small set of sensor nodes were able to determine that one of these scenarios was a strong match, as most sensors did not have a clear field of view of the evacuating group. While the accuracy of these matches was generally quite good with a typical correct:incorrect ratio of 4:1, in future work we hope to expand on the Predictive Simulation Reporting scheme to allow for local sensor nodes to collaborate together in isolating positive scenario matches. Local sharing of match confidence may reduce the number of incorrect matches without introducing substantial additional network traffic.

Our experiments featured a set of candidate scenarios which were pre-determined, featured a single crowd of 20 occupants, and contained at least one candidate that closely matched the proxy "real" evacuation progress. In future work we intend to implement the automated generation of future evacuation progress based on initial occupancy. As the number of possible combinations of routes occupants could take is very large, this motivates the development of an algorithm that generates a small number of candidates scenarios with high probability of one of the scenarios being a good match.

7 References

1. Chih-Yung Chen: A Fuzzy-Based Approach for Smart Building Evacuation Modeling. Innovative Computing, Information and Control (ICICIC). 2009 Fourth International Conference Issue Date: 7-9 Dec. 2009 pp. 1200 - 1203 (2009)
2. D. Dressler, M. Groß, J. Kappmeier, T Kelter, J. Kulbatzki, D. Plümpe, G.n Schlechter, M. Schmidt, Martin Skutella and S.Temme: On the use of network flow techniques for assigning evacuees to exits. Procedia Engineering Volume 3, 2010. pp. 205-215. 1st Conference on Evacuation Modeling and Management (2010)
3. S. Goel and T. Imielinski: Prediction-based monitoring in sensor networks: taking lessons from MPEG. ACM Comp. Commun. Rev. 31 5 (2001).
4. Y. Xu, J. Winter, W.-C. Lee: Dual prediction-based reporting for object tracking sensor networks. Proceedings of the 1st International Conference on Mobile and Ubiquitous Systems: Networking and Services (MobiQuitous'04). IEEE Computer Society, Los Alamitos, CA, USA, 2004 pp. 154–163. (2004)
5. Filippoupolitis, L. Hey, G. Loukas, E. Gelenbe, and S. Timotheou: Emergency Response Simulation Using Wireless Sensor Networks. The First International Conference on Ambient Media and Systems (Ambi-sys08). (2008)
6. Korhonen, T., Hostikka, S., Heliövaara, S., Ehtamo, H.,Matikainen, K.: Integration of an Agent Based Evacuation Simulation and the State-of-the-Art Fire Simulation. Proceedings of the 7th Asia-Oceania Symposium on Fire Science & Technology, pp.20-22 (2007)
7. R. M. Tavares and E.R. Galea: Evacuation modelling analysis within the operational research context: A combined approach for improving enclosure designs. Building and Environment 44 (5), pp. 1005–1016 (2009)
8. N. Pelechano , N. I. Badler: Modeling Crowd and Trained Leader Behavior during Building Evacuation. IEEE Computer Graphics and Applications, v.26 n.6, pp.80-86 (2006)