Performance of Signal Loss Maps for Wireless Ad hoc Networks

Henry Larkin, Zheng da Wu, Warren Toomey
Faculty of IT, Bond University, Australia
{hlarkin,zwu,wtoomey}@staff.bond.edu.au

Abstract. Wireless ad hoc networks face many challenges in routing, power management, and basic connectivity. Existing research has looked into using predicted node movement as a means to improve connectivity. While past research has focused on assuming wireless signals propagate in clear free loss space, our previous research has focused on using signal loss maps to improve predictions. This paper presents novel testing of signal loss maps in relation to the accuracy used for prediction purposes. Through analysis of test cases and results from a custom built simulator the performance is effectively measured.

Keywords. Signal loss maps, predicting signal loss, signal strength maps.

Introduction

In any wireless networking environment where nodes are utilizing predicted location information in order to improve routing, the ability to predict the communication strength from one location to another becomes vital. A prediction of the wireless connectivity between two nodes cannot accurately rely on location information alone. The majority of today's wireless networking environments feature many obstructions which reflect and block wireless signals. This rules out assuming all signals between two physical locations will travel with a clear line-of-sight approach, such as is relied on in [1]. In addition to needing knowledge about nodes' physical locations, the predicted propagation of signals over physical areas is required. Wireless transmission capabilities in an unknown or known environment are difficult to predict with accuracy [2]. Not only do different locations have different communication capabilities, but environmental conditions may change those capabilities over time. Foreign nodes operating on the same channel, radio-frequency interference, environmental noise and even landscape may change radically over time. This will affect any recorded or estimated measurements of signal loss.

To overcome these challenges, our previous work designed a signal loss map solution [3]. Signal loss maps represent the logical signal propagation topology over a physical area. They describe how signals are likely to propagate in various directions over various distances. Due to the constantly changing nature of the wireless environment, a perfect signal loss map is not possible to create with current technologies. However, various estimates may be developed to provide, with appropriate safety margins, predictions on whether two nodes at two locations will have connectivity in the future.

This signal loss map, dubbed the Communication Map (CM), is tailored to be built in real time using only wireless ad hoc nodes. The map is created using signal strength information provided with each packet as it is received from any node. To provide a physical reference system, some form of location-providing device is required for each node. In this research, a system such as GPS [4] is assumed to be available to provide the coordinates of each node. From these two external sources the CM is constructed.

The CM is made up of cells, defined areas which are square in shape and represent an average signal loss modifier. The signal loss modifier is a value which represents how a signal's loss increases over distance, relative to free space loss. This approach of using cells describes to users of the CM the same information that vendors of wireless cards use to describe range and signal strength capabilities. Vendors of wireless cards often include the maximum range and signal strength of their product in a variety of general scenarios, for example outdoors, home environment, cluttered office, etc. In a similar fashion, the CM of this research generates such scenarios in real time, and delineates where on a map such areas exist.

The general formula for calculating free space loss, S_{loss} (db), in ideal circumstances of wireless signals [5] is:

$$S_{loss} = 32.4 + 20Log_{10}F + 20Log_{10}D \tag{1}$$

where F is the frequency (MHz), and D is the distance (km) between the two nodes.

The free space loss formula may only be used *once* for the entire signal. To have an effect on the overall loss of a signal, each line *segment's* distance (how much a signal travels within a cell) will need to be adjusted by some modifier which represents a cell's effect on signal loss. To overcome this problem, the concept of *logical distance* is introduced. Any signal received or predicted using our CM is based on a *logical distance*. The *logical distance* that a signal travels is the distance it would need to physically travel in order to produce the same amount of loss, thus allowing the above formula to be used, given multiple signal modifiers.

Each cell represents the average signal loss of signals passing through that cell. The value stored for each cell is the modifier that a signal applies to the *physical distance* of a signal as it passes through that cell, which when multiplied together with the physical distance creates a *logical distance*. The minimum modifier value is 1.0, in other words a *logical distance* is identical to the *physical distance*, and thus represents perfect free-space loss. The modifier of each cell is used to extend the distance of a signal to the distance it would need to travel in perfect free space to achieve the same loss.

For an example of this process, consider Fig. 1 below. A signal will travel from Node A to Node B over the given CM. The direct line between the two nodes is formed, and a list of cells over which the signal will pass is created. Each of these cells multiply their modifier by the *physical distance* that the signal travels through their cell. The resulting *logical distance* can be used to find the signal loss, which can consequentially determine whether two nodes are predicted to be neighbours. This formula thus uses the modifier of each cell to *extend* the distance of a signal to the distance it would need to travel in perfect free space to achieve the same loss.

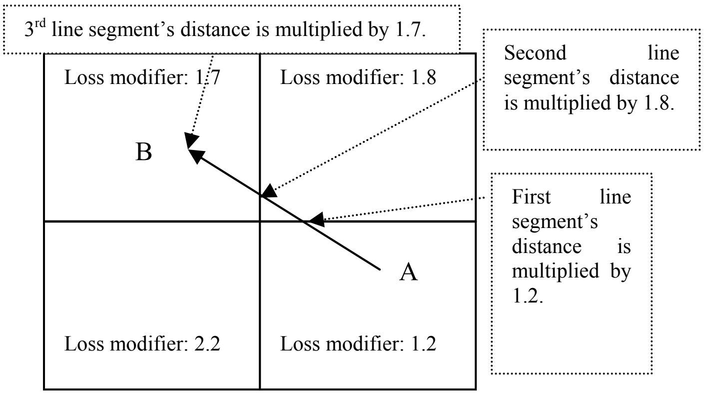

Fig. 1. Signal divided over multiple cells

Test Architecture

Existing simulators [6][7][8][9] already exist for basic network simulation. While many of these simulators are extensible, none specifically address the issues of wireless signal mapping and signal loss map testing. Because of this, a custom simulator was created. The simulator was created in response to fulfil the need for an appropriate testing bed for wireless ad hoc protocols that relied specifically on the needs of testing signal loss maps.

The simulator routes packets between nodes based on a simulated wireless environment map, which details user-created scenarios of how signals will propagate over real-world simulated objects. As each packet is successfully transmitted or lost (based on signal propagation), the simulator records what each node's CM predicted the signal loss to be, along with the actual signal loss based on the simulated wireless map. Fig. 2 below shows how a packet would be routed from a transmitting node to nodes within range, based on the simulated wireless environment map for the given scenario.

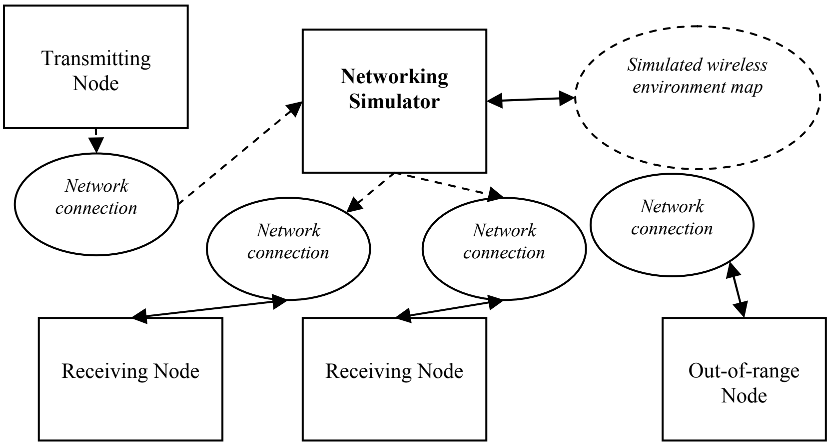

Fig. 2. Example Packet Broadcast

Scenarios

Several scenarios have been created to validate the concept of the Communication Map, while also identifying its weaknesses. The field of wireless communication has an unlimited number of practical scenarios. This presents a significant challenge in representing a broad spectrum of possible scenarios to gauge the overall effectiveness of the CM. A total of seven scenarios have been designed to both represent realistic environments and test the various aspects of the algorithms presented in this paper and in our previous work [3]. Some scenarios have been tested with a varying number of nodes to further analyse scenarios while also gaining an insight into the effects of network population.

Scenario 6 is presented as an example scenario. It contains only four nodes with two simulated wireless propagation areas. To the left is an area with a small signal loss modifier, such as that found in a forest. On the right-hand side is a building with reasonably thick walls (having a logical distance of 400 meters) and a reasonably high signal loss modifier of 5.0. All four nodes are in motion in this scenario, with nodes C and D moving from outside the building over to the forest area, and nodes A and B moving within the building. The scenario diagram is included in Fig. 3.

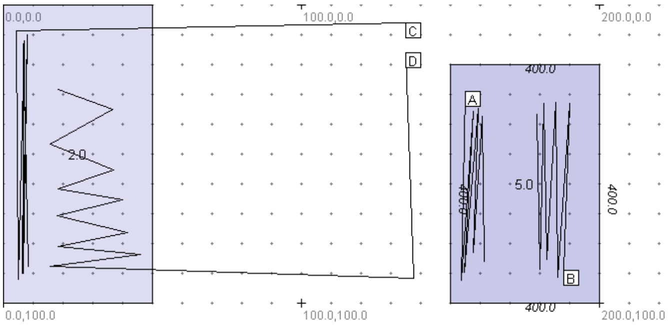

Fig. 3. Scenario 6 Overview

Results

Each of the seven scenarios were executed for each variation in settings and for each variation in node population where the scenario allowed multiple node populations. The results were then processed and summarised across different combinations of settings so that well-informed analysis could be conducted.

Boundaries and Default Cell Size (DCS)

The choice of implementation between boundaries and the basic CM is significant. The use of boundaries is a modification to the original algorithm as an attempt to provide possible improvement to the CM solution. With boundaries, each cell not only contains an average signal loss modifier for the area within the cell, but also additional modifiers for signals entering or leaving the cell through any of the 4 borders of the cell. The theory behind this was that certain real-world structures, such as buildings, had a far higher signal loss, but were not based over an area but rather a single object passed through between areas.

In the majority of these tests, two sets of results are generated for each scenario. One set is generated using the basic algorithm without boundaries, and one set is generated using the algorithm implementing boundaries. In each set, three Default Cell Size (DCS) settings are implemented, 25 meters, 50 meters, and 100 meters (the default). This is done to determine how reducing the DCS affects accuracy. In theory, the smaller the DCS, the more similar the results should be to a generic boundary implementation. This is due to the fact that a smaller DCS enables the Communication Map algorithms to more accurately map signal loss immediately. With larger DCS settings, the loss accounted for in artificial boundaries is averaged into the larger cells.

The first scenario tested for boundaries is the Scenario 4. This scenario was created specifically to test the boundary concept. The scenario consists of a single building, with perfect free-space loss within the building (for example in an empty warehouse) with basic walls of 150 logical meter loss (about that of a thin wooden wall). 3 nodes are placed within the building, with a further 5 nodes outside it. Two of the nodes, E and C, circumnavigate the building.

The results from these experiments are shown in Fig. 4 and Fig. 5. Without implementing boundaries, the average error stabilises between 35% and 50% using a DCS of 100 meters. Reducing the DCS by half improves accuracy considerably, with further reduction possible if the DCS is lowered to 25 meters. This demonstrates that a reduced DCS allows the CM to map some of the effects of simulated boundaries as implementing the boundaries algorithm does (when compared with Fig. 5). Implementing boundaries shows significantly improved accuracy, though more than this experiment alone is required to prove this.

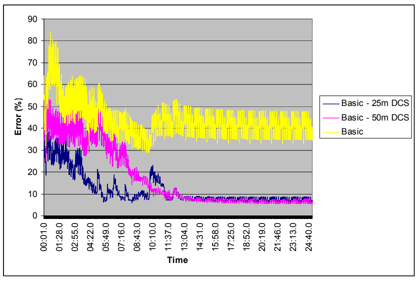

Fig. 4. Scenario 4 without Boundaries

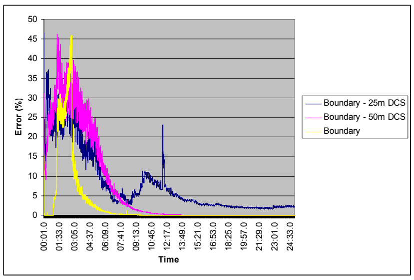

Fig. 5. Scenario 4 with Boundaries

Interesting to note that reducing the DCS does not improve accuracy in Scenario 4 with boundaries, but has the reverse affect. Fig. 6 - Fig. 9 illustrate the CMs at the end of each scenario from Node A's viewpoint. Fig. 6 shows how the CM algorithm represented the signal loss surrounding nodes A, B, and F by utilising the smaller cell size of subdividing the DCS as a boundary representation itself. Fig. 7 shows the typical 100m DCS with boundaries, with Fig. 8 and Fig. 9 reducing the DCS to 50m and 25m respectively. The smaller DCS with boundaries hinders the boundary development, as boundaries are more difficult to develop than cells. Fig. 6 without boundaries shows that even with 8 nodes the CM does not perfectly represent even the shape of the simulated building, let alone the accurate signal loss. Without a thorough grid of cells there will be varying quantities of signals mapped over various cells, simply from node positioning and node movement. A perfectly accurate CM is impossible to develop using current technology. With the smaller DCS on the boundary examples, there become too many objects where signals can be mapped. As signals being mapped are averaged over all objects along the assumed signal's path, the greater number of objects requires a greater quantity of node movement to further correlate signal loss to more accurately-placed objects.

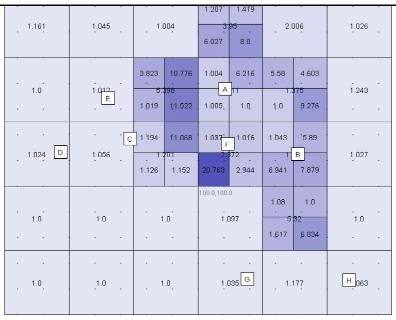

Fig. 6. Scenario 4: 25m DCS without Boundaries

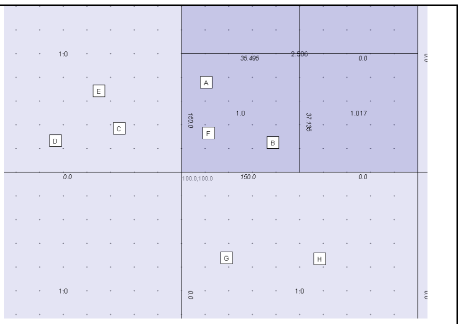

Fig. 7. Scenario 4: 100m DCS with Boundaries

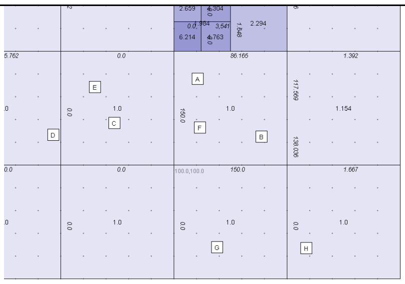

Fig. 8. Scenario 4: 50m DCS with Boundaries

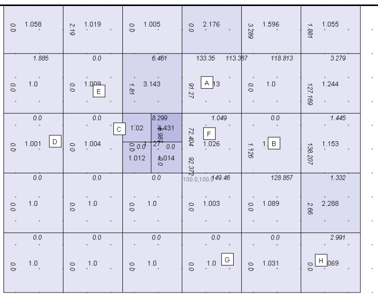

Fig. 9. Scenario 4: 25m DCS with Boundaries

To further see the effects of implementing boundaries more scenarios are needed. The same experiments were run on Scenario 5 and Scenario 6, both of which also have high boundary use yet simple simulated environments that have a better chance of being mapped. Results are shown in Fig. 10 and Fig. 11. These experiments indicate that overall the varying DCS settings and the implementation of boundaries make little difference to the overall accuracy.

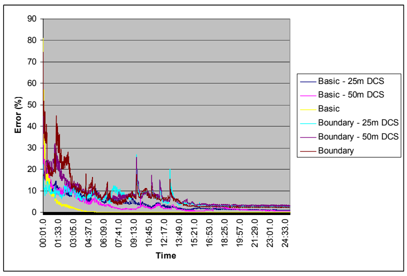

Fig. 10. Scenario 5 Results

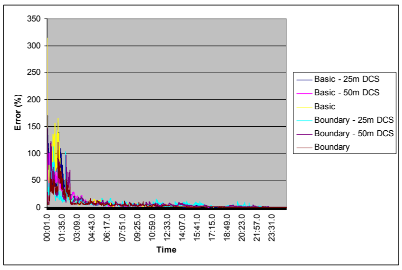

Fig. 11. Scenario 6 Results

The experiments were then run on Scenario 7, which is based on a real street. Despite the number of simulated objects and boundaries, all settings produced similar results (refer to Fig. 12 and Fig. 13). The CMs at the end of the simulations of a 25m DCS without boundaries and a 100m DCS with boundaries experiment are shown in Fig. 14 and Fig. 15 for interest. The favourable results in these experiments can be attributed to the low amount of node movement, where only three nodes are moving in relatively fixed movement patterns. A perfect CM is not required to produce favourable results, only a CM which adequately represents the signal loss as it has and will be used in the future.

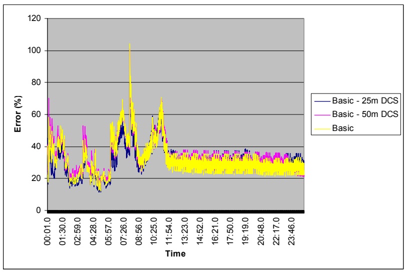

Fig. 12. Scenario 7 without Boundaries

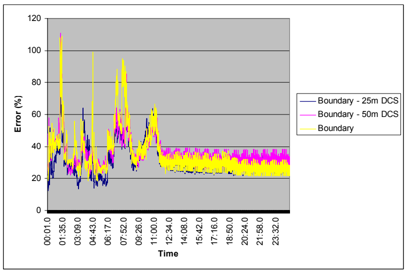

Fig. 13. Scenario 7 with Boundaries

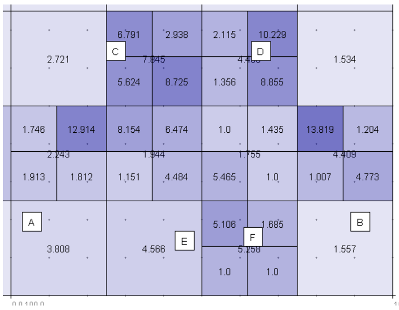

Fig. 14. 25m DCS CM without Boundaries

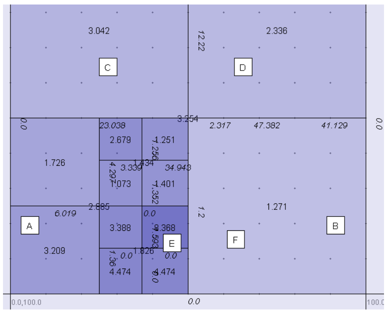

Fig. 15. 100m DCS CM with Boundaries

Number of Nodes

Another area of interest was to determine if the number of nodes had an effect of the accuracy of the CM. Fig. 16 and Fig. 17 graph all setting samples over Scenario 1 and Scenario 2 respectively. These are the two main scenarios where the basic layout would allow the number of nodes to play an influential role without new areas being discovered by the increased number of nodes. Scenarios 3 and 4 were included (Fig. 18 and Fig. 19) were tested also, though these examples focus mainly on boundary testing.

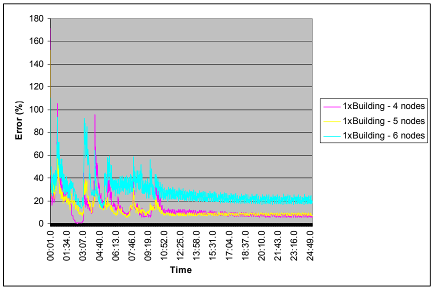

Figure 16. Scenario 1

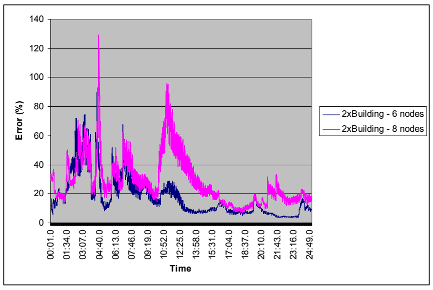

Figure 17. Scenario 2

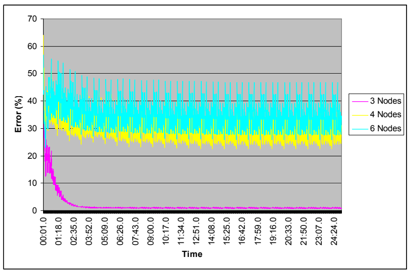

Fig. 18. Scenario 3

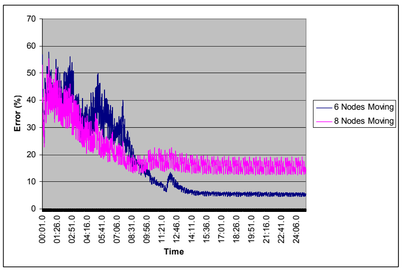

Fig. 19. Scenario 4

The results in these experiments were surprisingly different than anticipated. In theory, the number of nodes would have favourably increased the accuracy of the CM. However, in almost all scenarios tested, an increase in the number of nodes had an adverse effect on accuracy. In Fig. 16, using both 4 and 5 nodes resulted in very similar results, with 5 nodes performing slightly better, as would be expected. Using 6 nodes, however, almost doubled the overall average error.

In Scenario 2 (Fig. 17), 8 moving nodes generate a more accurate map than 6 moving nodes, but only within the first 10 minutes of the experiment. Towards the end of the simulation the average error increases, and unfortunately stabilises this way due to the lack of future node movement. This is simply a case of a new area of signal loss being discovered, where a lack of node movement fails to accurately identify that loss. In Scenario 3 (Fig. 18), the best performance is surprisingly from using only 3 nodes. While the actual CM built using only 3 nodes is far from being an accurate representation of the simulated signal loss, it provides enough detail given the information the nodes have accumulated.

Both in theory and in the experiments performed, a greater number of nodes has a greater chance of producing a more accurate CM. However, actual prediction accuracy depends largely on how the CM is used by nodes during the scenarios. Nodes which remain in the same approximate area both further improve the accuracy of that area and benefit from the effects. Nodes which use the CM in untested areas are less likely to obtain ideal results. Due to the vast nature of exact node movement and positioning, various experiments will achieve various results. Overall, however, a more accurate CM is created given a greater number of nodes.

MCSD and MCMD

The Minimum Cell Subdivide Difference (MCSD) and Minimum Cell Merge Difference (MCMD) values have also been considered. These values govern when cells are able to subdivide (split into 4 equal parts) and merge back together when the level of detail required by the CM changes over time. By default, all scenarios use a MCSD of 3.0 and a MCMD of 2.0. This means that if the average modifier within a cell varies by more than 3.0 then the cell will subdivide to allow for a greater level of detail to be mapped. If the average modifier between 4 adjacent nodes falls to less than 2.0, then the cells are merged back together again. These settings have been modified in two groups of settings; the first uses a MCSD of 2.0 and a MCMD of 1.0, and the second uses a MCSD of 1.0 and a MCMD of 0.0. These are all compared with basic settings in Fig. 20 below.

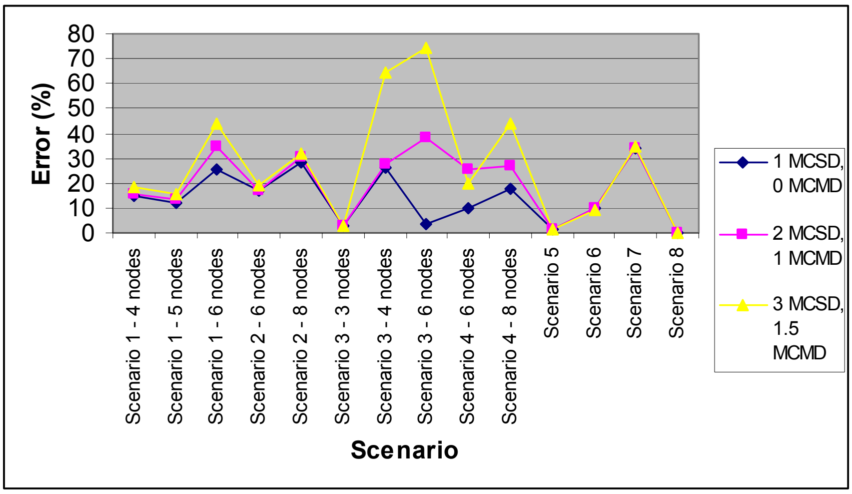

Figure 20. MCSD and MCMD Performance

The results from these experiments show that lowering the MCSD and MCMD values increases accuracy. The reasoning behind this is that lower MCSD and MCMD values allow cells to subdivide faster and merge back together without difficulty. As previously discussed with DCS values, the smaller the cells the more accuracy is obtained, and lowering the MCSD and MCMD values have the same effect. However, a reduced DCS value increases bandwidth costs.

The scenarios focusing specifically on boundary-optimised situations benefit the most from lower MCSD and MCMD values, as smaller cells more closely represent boundaries. The most interesting results are that the more realistic scenarios (Scenario 2, Scenario 6, and Scenario 7) show almost no difference in changing the MCSD and MCMD. From this it is concluded that while the MCSD and MCMD values in theory have an effect on accuracy, in practice the Communication Map algorithms perform well regardless of these settings.

Conclusion

Predicting connectivity between nodes based on location information can be improved with an understanding of wireless signal propagation in each environment. Signal loss maps generate this information, and our Communication Map solution achieves this in real time without user intervention. The use of signal loss maps in wireless ad hoc routing is a relatively new field, and one which requires significant performance analysis before the advantages become evident. This paper has presented such testing under a number of custom-created scenarios, as well as the various settings which may be used to improve accuracy. The results of these tests conclude that signal loss maps can perform accurately under a variety of situations. Further testing can be conducted if the CM is applied to existing routing protocols, and performance investigated.

References

1. Su, W. W., Motion Prediction in Mobile/Wireless Networks. PhD dissertation. University of California, Los Angeles, USA. 2000.
2. Howard, A., Siddiqi, S., Sukatme, G. S., An Experimental Study of Localization Using Wireless Ethernet, in 4th International Conference on Field and Service Robotics, July 2003.
3. Larkin, H., Wireless Signal Strength Topology Maps in Mobile Adhoc Networks, Embedded and Ubiquitous Computing, International Conference EUC2004, pp. 538-547, Japan. 2004.
4. Enge, P., Misra, P. Special issue on GPS: The Global Positioning System. Proceedings of the IEEE, pages 3–172, January 1999.
5. Shankar, P. M., Introduction to Wireless Systems. Wiley, USA. 2001.
6. Fall, K. Network Emulation in the Vint/NS Simulator. Proceedings of ISCC'99, Egypt, 1999.
7. McDonald, C.S., A Network Specification Language and Execution Environment for Undergraduate Teaching. Proceedings of the ACM Computer Science Education Technical Symposium '91, San Antonio, pp25-34, Texas, Mar 1991.
8. Unger, B., Arlitt, M., et. al., ATM-TN System Design. WurcNet Inc. Technical Report, September 1994.
9. Keshav, S. REAL: A Network Simulator, tech. report 88/472, University of California, Berkeley, 1988.