

Information Dissemination in Mobile Ad-Hoc Geosensor Networks

Silvia Nittel¹, Matt Duckham², and Lars Kulik³

¹ National Center for Geographic Information and Analysis
University of Maine, Orono, ME 04469, USA

nittel@spatial.maine.edu

² Department of Geomatics

University of Melbourne, Victoria, 3010, Australia

mduckham@unimelb.edu.au

³ Department of Computer Science and Software Engineering
University of Melbourne, Victoria, 3010, Australia

lkulik@cs.mu.oz.au

Abstract. This paper addresses the issue of how to disseminate relevant information to mobile agents within a geosensor network. Conventional mobile and location-aware systems are founded on a centralized model of information systems, typified by the client-server model used for most location-based services. However, in this paper we argue that a decentralized approach offers several key advantages over a centralized model, including robustness and scalability. We present an environment for simulating information dissemination strategies in mobile ad-hoc geosensor networks. We propose several strategies for scalable, peer-to-peer information exchange, and evaluate their performance with regard to their ability to distribute relevant information to agents and minimize redundancy.

1 Introduction

Increasing decentralization is a widespread feature of information system architectures, made possible by the advances in computer networks over the past two decades. Decentralized information systems are acknowledged as offering several advantages over centralized architectures, including improved reliability, scalability, and performance [1].

In the context of mobile and location-aware systems, centralization remains the dominant system architecture today. For example, location-based services, which aim to provide more relevant information to mobile users based on information about their geographic location, typically adopt a centralized architecture [2]. In a conventional location-based service, each mobile user accesses information services from remote service providers, which perform the task of capturing, managing, and updating any information relevant to their application domain.

In such a system, the centralized remote service provider can act as a weak point in the system. The bottleneck of a single access point decreases system

reliability and performance. The system is not scalable, because additional load from new users and services must be borne by the service provider. With respect to the domain of geographic information services, the service provider will also face the problem of regular maintenance and update of complex spatiotemporal data, such as environmental conditions or traffic events.

One topic where decentralized architectures are already a vital feature is *geosensor networks*. Advances in sensor technology and the development of inexpensive small-form, general-purpose computing platforms have led to the study of *sensor networks*. Sensor networks comprise multiple miniature “PCs”, each of which contains a CPU, volatile and stable memory, short-range wireless communication, battery power, and attached sensors. The on-board sensors are used to collect information about the physical world, like temperature, humidity, or the current location of objects. Sensor nodes can be deployed in high density within the physical world and enable the continuous measurement of phenomena in unprecedented detail. A geosensor network is defined as a sensor network that monitors phenomena in geographic space [3].

In this paper we apply the paradigm of geosensor networks to the problem of decentralized location-based services. Geosensor networks rely on decentralized architectures, but their sensor nodes are currently rarely mobile and primarily concerned with information capture rather than information service provision. By contrast, location-based services do provide information services to mobile users, but rely on a centralized architecture. The key question facing a decentralized location-based service is how to disseminate relevant geospatial information to spatially dispersed mobile users. The core focus of this paper is to examine this question in more detail.

In Section 2 we cover the related work and the background to the problem, with a particular focus on geosensor networks and mobile ad-hoc networks. In Section 3, we discuss the problem setting and motivation in more detail. Section 4 presents three related strategies for geospatial information dissemination in a geosensor network, and describes a simulation environment developed to enable the performance of different information dissemination strategies to be tested. Section 5 presents the initial results of using the simulation environment. We conclude with suggestions for further research in Section 6.

2 Background

The continuing development of ubiquitous wireless communication technology, including miniaturization of computing platforms and the development of nano-scale sensors, is enabling new computer applications that would not have been possible in the past. Location-based services are one of the most recent, influential application domains; users with hand-held devices employing location-sensing and wireless communication are able to retrieve up-to-date information related to their immediate geographic environment (see [4] for a review of location-sensing techniques). In this section we review the relevant literature within two fields: geosensor networks and mobile ad-hoc networks.

2.1 Geosensor Networks

Recent and projected advances in small, low-cost microelectronic and mechanical systems (MEMS) with limited on-board processing and short-range wireless communication capabilities are also changing the way that we collect and process information about the physical world [5, 6]. Small, inexpensive sensors can be attached to physical objects or embedded in the environment, and sense as well as process the information that is collected. Today, networked sensors can be constructed by using open source and commercial components at the size of an inch or smaller such as the Berkeley/Intel Mica Motes [7, 8]. Other examples of powerful sensor network platforms are the MIT Cricket [9] and the UCLA WINS systems [10].

Large collections of untethered, battery powered sensors with various sensing functions can be distributed over a geographic area, and measure traffic and road conditions, environmental indicators, or seismic activity at a fine-grained temporal and spatial scale that was not possible in the past [11–13]. Since such sensor nodes are tiny and the limited battery capabilities allow only for short range wireless communication, they must communicate with peer sensor nodes in their spatial proximity. Projecting the continued miniaturization, it is not expected that sensor nodes connect to a centralized computing server to upload or stream data directly; they might, however, communicate with a local “base station,” i.e. a more powerful sensor node with larger processing, storage, communication and energy capabilities. In general, information is routed via multiple network “hops” to a centralized server and applications [14–16], or the information is processed in the local geographic environment of sensor node.

Integrating both location-aware systems and sensor network technology, we can envision sensor nodes that are aware of their geographic location, equipped with diverse sensors, mobile, and can communicate with nodes in close spatial proximity concerning information that they have collected. Sensor nodes can be mobile by either being self-propelled or by being attached to moving objects, like automobiles, USPS packages, or even humans. As defined above, a geosensor network is a network of sensor nodes that are aware of their geographic location, and capture information about environmental phenomena via on-board sensors [3].

2.2 Mobile Ad-Hoc Geosensor Networks

Efficient information routing is a significant research challenge in geosensor networks, and sensor networks at large [14, 17]. Preserving battery life is the key design criteria for designing communication protocols in order to maximize system lifetime once a sensor network is deployed. Communication is a significantly higher drain on the battery than processing information locally on a sensor node. Furthermore, once nano-sensors are deployed in the environment the battery packs cannot be replaced.

In contrast to information routing in many of today’s communication networks, which is address-based (IP addresses), routing in sensor networks is data-centric. The goal is to distribute information only to sensor nodes that need the

information or that can be a source of information [16, 18]. Another aspect of data-centric routing is scalability: if the number of nodes in a sensor network increases to thousands or millions of nodes, a decentralized, peer-to-peer information dissemination and data collection strategy can provide efficient information distribution and eliminate the bottlenecks of a centralized database or service architecture.

The geosensor networks envisioned in this section can be seen as a type of *mobile ad-hoc network* (MANET, a self-configuring wireless network of autonomous mobile nodes). To emphasize this connection, in the remainder of this paper we refer to these geosensor networks as MAGNETs: *mobile ad-hoc geosensor networks*. Efficient location-based information dissemination strategies are highly relevant for MAGNETs [19]. Geosensor nodes in a MAGNET capture information that is relevant in a geographically constrained context, i.e., in close proximity to the event (e.g., a hazard warning). The location-dependency of information in a MAGNET contrasts strongly with generalized computer networks, such as the Internet, in which the storage location of information may be entirely independent of any locations to which that information refers.

Since sensor nodes often leave the neighborhood of an event, an important research question is how to achieve efficient sharing of relevant information between sensor nodes. This question has already begun to be answered by [20, 21]. The information dissemination strategy explored in this paper, where agents communicate with one another whenever they happen to be in close spatiotemporal proximity, is similar to the *opportunistic exchange* described in [21, 22].

Traditional MANET information routing strategies with the objective to route information from a specific source to a destination can be categorized into two classes: *table-driven* and *on-demand*. In table-driven routing, each node proactively maintains up-to-date information about the routing paths between every pair of nodes in the MANET. In on-demand routing, a source node reactively generates a route to a destination node when required, based on responses to a query that floods the network. Table-driving routing strategies are generally thought to be more suitable for larger MANETs with high levels of mobility [23, 24] (e.g., MAGNETs) although this result has been called into question by some other studies [25]. Opportunistic exchange is orthogonal to the MANET routing strategy, and may be used in addition to, or independently of, table-driven or on-demand routing. In this paper we assume a “routing-free” model, where information dissemination is purely opportunistic, and geosensor nodes do not need to explicitly query other nodes for information. Instead, the spatiotemporal location of a geosensor node can be thought of as an *implicit query* for any information that concerns locations in close spatiotemporal proximity.

3 Problem Statement

The previous sections have provided a background to geosensor networks and argued that collaboration strategies of geosensor networks represent a new paradigm for location-based services. For example, consider a conventional

location-aware navigation system, where a centralized location-based service provider stores and manages real-time information about traffic congestion. Drivers accessing the service might expect to receive continuously updated information from a centralized service provider concerning the locations of traffic jams. In order to provide this service, the service provider would need to address all the problems of scalability, reliability, and performance discussed previously. Computing and communicating relevant customized information for each LBS client will be a significant performance problem for a centralized service with an increasing number of clients.

However, by adopting the geosensor networks paradigm, each vehicle may be thought of as a mobile geosensor node, able to sense information about its own location and local traffic conditions, process this information, and communicate this information to other vehicles in its neighborhood. Potentially, this decentralized location-based service could offer improved reliability and performance, because there exists no centralized service provider acting as a bottleneck to information dissemination and processing. Scalability issues are also positively affected, since rather than increasing the pressure on a single service provider, adding new vehicles to the system increases the number of available nodes for information dissemination and processing. As we shall see in later section, increasing the nodes within a decentralized location-based service can actively *improve* rather than degrade system performance.

In addition to the general advantages of decentralized information system architectures (scalability, reliability, and performance), we can identify at least two further advantages that are peculiar to the domain of location-based services:

1. Many types of geographic information, such as information about traffic conditions and weather, are highly *volatile* in the sense that they can change rapidly, both spatially and temporally. Using a centralized architecture, where all updates must be processed through the service-provider bottleneck, makes it harder for a service to respond rapidly to changes in volatile information.
2. Communicating rapidly changing spatial and temporal information to a centralized service provider may introduce considerable redundancy into the system. Information that refers to a very specific location and time will not be relevant to the vast majority of service users. Therefore, processing such information using the same channels as information with wide spatiotemporal relevance is not an efficient use of limited communication and computational resources.

Continuing the example of the location-aware navigation system, changes in traffic congestion or road conditions will be continually detected by vehicles moving through the environment. If all such updates must be submitted to a centralized service provider, this will further increase the performance bottleneck of the system. For rapidly evolving phenomena, such as traffic congestion or taxi cab passengers looking for a ride, information may become outdated fast, possibly before the centralized service provider is able to make updated information

available to other vehicles. Further the information may only be relevant to a small percentage of the overall service users, yet it may require as much centralized communication, processing, and storage resources as information that is vital to many or even all service users.

The example of a decentralized traffic information system used above is not so futuristic. Location-aware sensors are common in modern vehicles, as are environmental sensors able to track environmental conditions (like temperature, humidity, or light conditions). A vehicle may even be able to derive road conditions from its on-board sensors such as windshield wipers, brakes, speed, and so forth [26]. Furthermore, the number of such sensor-enabled devices in the environment is steadily increasing as the technology costs decrease. With hundreds, thousands, or even millions of such sensor-enabled devices in the environment, the boundary between location-based services and geosensor networks is set to blur.

Conversely, from the perspective of geosensor networks, there is increasing interest in sensor node mobility. For example, a key application area for geosensor networks is microclimate monitoring. Hundreds or thousands of sensors distributed throughout a region can continuously sense environmental conditions, for example, temperature and humidity throughout a wine-growing region or for precision agriculture. In the future it may also be possible for such geosensor networks to reconfigure, with sensor nodes moving to the optimal location for responding to particular user queries.

A major issue facing these decentralized, ad-hoc, geosensor networks is the formulation of an efficient *information dissemination strategy*. Different information dissemination strategies vary in the way they are able to address questions such as: (1) If new information becomes available, how long does it take before the network stabilizes? (2) What is the size of the optimal distribution radius for the information origination? (3) How long does the sensed information persist in the network and what parameters can ensure persistence of information?

In the following section we distinguish between three different strategies. *Flooding* is where each geosensor node that encounters an event or receives a message about an event passes on the information to every other node within its communication range. The second approach is referred to as an *epidemic*, in which each node only informs n other nodes about events. The third approach is *location-constrained*, in which information is only passed on in proximity to the event, and then discarded.

4 Approach and Simulation Model

In the discussion above, we argued that it will be an advantage for the next generation of location-based services to adopt a decentralized architecture, akin to geosensor networks. In this section we begin to explore the precise nature of efficient information dissemination strategies based on localized communication between agents in a geosensor network. In this context, an agent is defined as an autonomous system that is situated within an environment, senses its environ-

ment, and acts on its sensed knowledge of its environment [27]. Specifically, we are concerned with:

mobile location-aware agents, able to sense information about their immediate geospatial environment and communicate with other agents in their neighborhood.

The goal of a MAGNET is to ensure the *efficient* communication of *relevant* information between such agents. In this context, relevance refers to the pertinence of information to the task or tasks in which an agent is engaged [28]. Efficiency (or more precisely inefficiency) can be decomposed into two key features: *ignorance* and *redundancy*. Ignorance concerns the situation where an agent fails to receive relevant information in time. Redundancy concerns the situation where an agent successfully receives irrelevant information. We will return to issues of relevance and efficiency in the following section.

To begin to investigate the potential of MAGNETs, we have developed a prototype simulation testbed, using Java. The reason for favoring a simulation approach, at least initially, is that the wide variety of possible application domains and geosensor network configurations currently precludes a more analytical approach. The prototype is intended to allow researchers to gain insight into the possible effects of manipulating different parameters and their effect on information dissemination strategies in geosensor networks.

The simulation testbed enables the manipulation of six broad classes of parameters, outlined below:

- *Environment*: A variety of different environments can be used in the testbed. The environments can be simulated or derived from information about real geographic environments. Currently, only one- and two-dimensional spatial environments (e.g. graphs and the Cartesian plane, respectively) are supported by the software, although support for three-dimensional dynamic environments would be possible with limited modifications.
- *Communication*: Limited communication capability is a fundamental feature of agents in geosensor networks, which constrains an agent’s access to information. Limitations on communication range, frequency, and latency can all be modelled and manipulated within the simulation testbed. Further, different communication protocols can be implemented, discussed in more detail in the example simulation that follows.
- *Sensors*: The variety of sensors available to an agent determines the types of information that can be sensed. Other characteristics that may be modelled that include the sensor reliability, accuracy, and sensor granularity.
- *Mobility*: An agent’s mobility characteristics limit the environments and neighborhoods an agent can access, and so the opportunities for sensing and communication. Mobility parameters such as speed of movement, patterns of movement (e.g., goal-directed movement or random walks), and constraints to movement (e.g., by the environment or agent congestion) are available simulation parameters.

- *Tasks*: The task an agent is performing determines whether received or sensed information is relevant to the agent. Further, an agent’s task will affect other aspects of an agents behavior, such as patterns of mobility. It can also influence other agents in the system.
- *Agents*: In addition to agent mobility and tasks, discussed above, other characteristics of the agents that may be manipulated using the testbed include agent memory, agent information processing capabilities (e.g., spatial analysis or interpolation), and agent life cycle (e.g., when and how agents enter and exit the simulation). The constraints of limited agent memory have already been explored in [21].

The parameters described above may be varied both spatially and temporally. For example, an agent’s speed or patterns of mobility can vary spatiotemporally, e.g., modelling a car caught in city center traffic congestion during rush hour. In the following section we examine some initial results of using the simulation testbed.

5 Example Simulation

The wealth of parameters that can affect a MAGNET leads to a large potential solution space for optimal information dissemination, with parameters that may vary according to the specific application domain. Given limited communication resources (primarily bandwidth and communication range) a vital question facing the agents in a geosensor network is “Under what circumstances should an agent transmit information to another agent in its neighborhood?” We classify three distinct classes of communication strategy in the following way:

1. *Flooding*: an agent always informs all other agents within its communication range of all information about “events” it has collected (either sensed or received from other peers).
2. *Epidemic*: an agent informs only the first n agents it encounters within its communication neighborhood after discovering an event.
3. *Proximity*: an agent informs other peers within its communication neighborhood only as long as the agent is within a certain threshold distance d from the event location.

5.1 Traffic Hazard Warning Simulation

To make our proposed approach more concrete, we consider a traffic hazard warning system based on a MAGNET. We assume road users (mobile agents) are driving around in a road network (environment) in location-aware vehicles equipped with sensors that are able to detect hazards (e.g., the presence of icy road conditions, road work, accidents, or traffic congestion). We also assume vehicles possess short-range peer-to-peer communication capabilities. Although battery power is clearly not an issue for geosensors attached to a significant power source like a vehicle, communication capabilities will in general still be

limited, for example, by bandwidth. Consequently, the general form of our problem statement, that agents can only communicate with other agents in their immediate neighborhood, remains valid.

The goal of the geosensor network is then to enable relevant information about road hazards to be efficiently communicated to agents as they travel on their journey. In order to simulate this application we make several simplifying assumptions (which we will begin to relax later). First, we assume a simulation environment comprising a fixed regular grid of connected roads and a fixed single point-location hazard located fairly centrally within the grid. Second, we assume fixed and constant speed of movement as well as communication range for all vehicles. Third, we assume road users are engaged in goal directed movement: users begin and end their journeys at randomly selected locations within the environment, but take the shortest path between those locations. Fourth, the number of agents within the simulation is fixed and constant: as one agent ends its journey and leaves the simulation, another begins its journey and enters the simulation.

Figure 1 shows four of the initial iterations for this simplified simulation. The simulation is discrete, so at each iteration agents move a fixed distance along the shortest path toward their destination. Figure 1 depicts a part of the environment (the road network shown as a gray grid), the location of a simulated hazard (indicated with an exclamation mark), and the locations of agents that know about the hazard (black dots). Agents discover information about the hazard either by sensing it directly or by being informed about the hazard by other agents. The spread of information about the hazard shown in Figure 1 was produced using the flooding communication strategy.

5.2 Simulation Results

In addition qualitative simulation maps (e.g., as shown in Figure 1), we used the simulation testbed to generate quantitative information about the performance of the simulated information dissemination strategies in the geosensor network.

Ignorance. First, we tested the levels of ignorance in the different chosen information dissemination approaches. Figure 2 depicts the levels of ignorance within the system over the first 200 iterations following the initial discovery of the hazard (averaged over 5 independent simulation runs). For this simulation, 200 iterations were sufficient for the system to stabilize. In Figure 2, agent ignorance i is measured as:

$$i = 100. \left(1 - \frac{|E \cap K|}{|E|} \right)$$

where E is the set of agents who encounter the hazard and $E \cap K$ is the set of agents who encounter the hazard knowing about it in advance. Thus, the ignorance measure i varies as a percentage from 0 (total ignorance, no agents

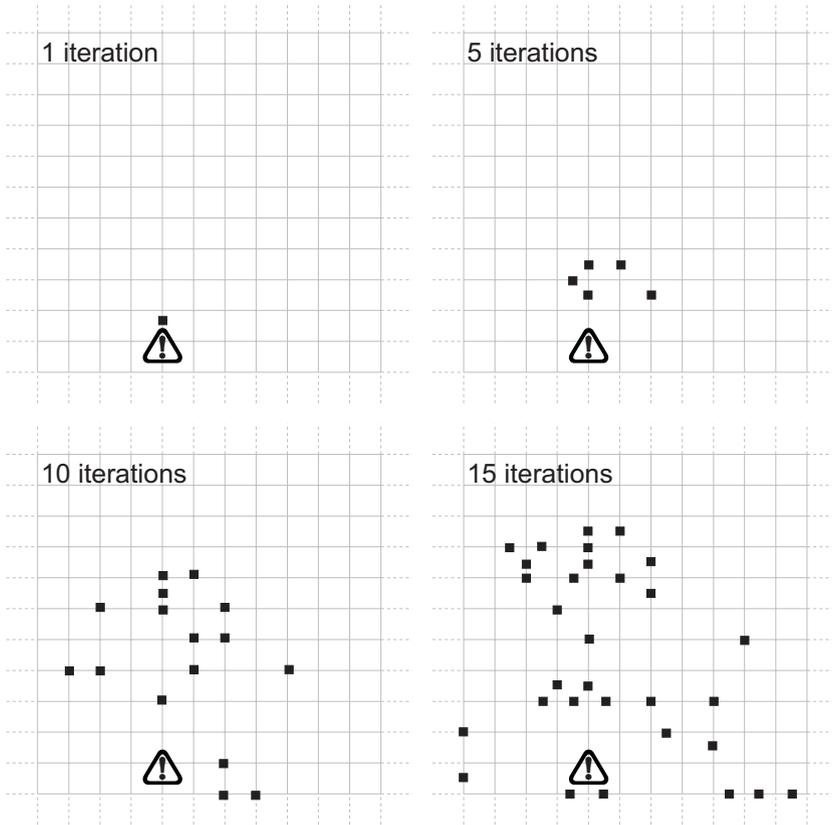


Fig. 1. The spread of information: Location of agents which know about the hazard after 1, 5, 10, and 15 iterations (using the flooding communication strategy for part of the environment)

who encountered the hazard knew about it in advance) to 100 (total knowledge, all agents who encountered the hazard knew about it in advance).

To produce Figure 2 we used a fixed number of 100 agents in the system, 736 road network nodes, and a communication range for an agent of 10 nodes, on average. Although the communication distance was fixed, the exact number of nodes within communication range varies slightly depending on the position of the agent in the network (e.g., fewer accessible nodes near the edge of the road network). Together the number of agents, environment size, and communication range can be combined to yield the probability that at any given point in the simulation an individual agent will be within communication range of at least one other agent (discussed in more detail later).

Figure 2 shows the differences in information dissemination performance across the three different communication strategies. In terms of agent ignorance, the flooding communication strategy clearly offers the best and most robust

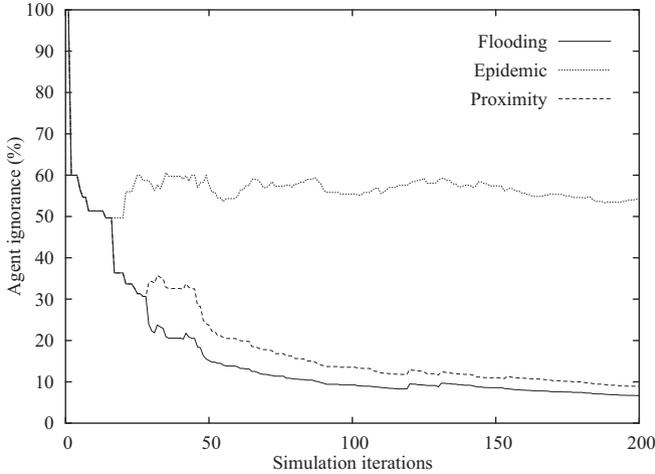


Fig. 2. Agent ignorance measure

performance: more than 90% of agents acquire the knowledge they need by the 200th iteration. The proximity strategy performs almost as well, while the epidemic strategy initially performs well, but reaches a performance ceiling at about 50% agent ignorance, i.e. 50% of the vehicles that encounter the hazard had not received the warning information from other agents ahead of time. As mentioned previously, in all our initial simulations hazards are fixed and static.

Redundancy. As discussed above, redundancy is another performance parameter in geosensor network efficiency, especially with regard to system scalability. In Figure 3, redundancy r is measured as:

$$r = 100 \cdot \left(1 - \frac{|E \cap K|}{|K|} \right)$$

where K is the set of agents who know about the hazard and $E \cap K$ is the set of agents who encounter the hazard knowing about it in advance (as above). The figure shows high levels of redundancy across all communication strategies, although the proximity strategy equilibrates at slightly lower levels of redundancy than the other two communication strategies.

Cost of Redundant Messages. Based on the evidence presented so far, the flooding strategy appears to be the primary choice for a dissemination strategy, since it achieves the lowest level of ignorance and a level of redundancy marginally worse than the proximity strategy. These parameters are initial guides, but not the only possible measures of efficiency with regard to geosensor networks. Since a large percentage of agents in geosensor networks operate in a battery- and communication bandwidth-restricted environment, the cost of

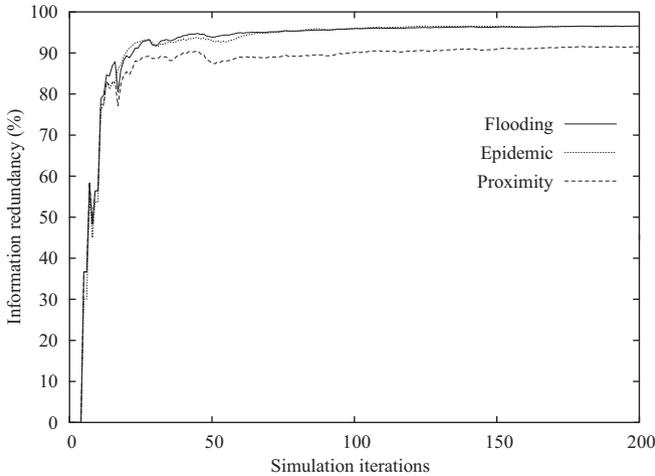


Fig. 3. Information redundancy

sending redundant messages to inform an agent about an event is a significant performance measure for a proposed information dissemination strategy. For instance, it is sufficient that an agent knows about a hazard, and is informed by one other peer (best case). However, some strategies such as flooding produce the effect that an agents “hears” of a hazard over and over again. An alternative measure of redundancy, then, is the *total number* of messages sent by agents using each communication strategy (Figure 4). These results show that the flooding communication strategy produces a significantly larger amount of messages than the proximity strategy, which in turn produces more messages than the epidemic strategy. Thus, robustness is achieved at the cost of a high degree of message passing overhead. After equilibration (following the first 50 iterations) the number of messages increases roughly linearly, suggesting that the rate of message increase could be another useful redundancy index for simulations in a steady-state.

Other Parameter Variations. Another factor that strongly affects efficiency is the degree of peer-to-peer connectivity. The number of agents, the size of the environment, and the communication range of agents can be combined to yield the probability that, at any particular iteration, an arbitrary agent is within communication range of at least one other agent. For simplicity, we refer to this probability in the following text simply as the *probability of communication*, or $P(C)$. The probability of communication is the preferred measure of connectivity, since it provides a measure of peer-to-peer connectivity that is independent of the specific context of the simulation. For a population, a , of 100 agents, an environment size, e , of 736 nodes and a connection range, c , of on average 10 nodes, the probability of communication can be computed as approximately 0.75 using the following formula:

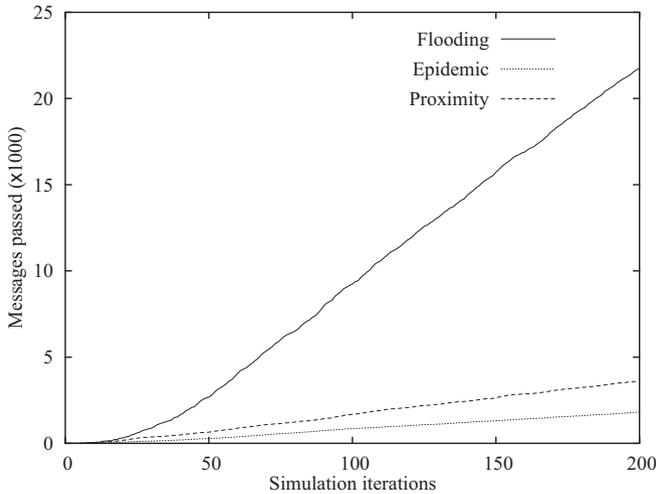


Fig. 4. Messages sent over initial 200 iterations

$$P(C) = 1 - \left(\frac{e - c}{e} \right)^a$$

Within the simulation testbed, connectivity can be investigated by varying the probability of communication and measuring the steady-state levels of agent ignorance in the system. Figure 5 shows the steady-state levels of agent ignorance achieved by the three different communication strategies at different levels of peer-to-peer connectivity. The probability of communication was calculated using the formula above, by varying the number of agents in the simulation, and maintaining a constant environment size and communication range. The agent ignorance level shown in Figure 5 is the average over 100 iterations of 5 simulation runs after each simulation has equilibrated. The error bars in Figure 5 show the variability of one standard deviation for this population of observations.

The figure shows that the proximity strategy is the most sensitive to changes in connectivity, performing almost as well as the flooding strategy at the highest levels of connectivity. The figure also indicates that higher levels of connectivity generally lead to lower variability in ignorance. This in turn may be taken to indicate that higher levels of connectivity lead to more stable geosensor networks, where performance is liable to vary less.

5.3 Summary

Overall, the proximity communication strategy seems to provide a favorable compromise in terms of MAGNET efficiency. The levels of agent ignorance achieved using the proximity strategy are comparable to those of the flooding strategy. At the same time, the proximity strategy does not lead to as high levels of information redundancy as the flooding strategy.

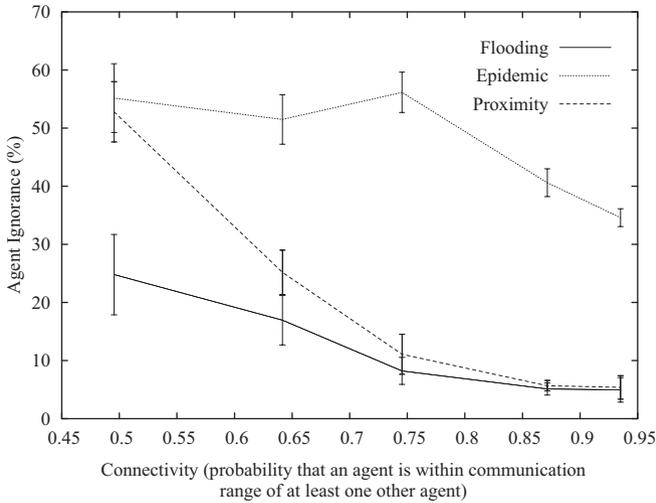


Fig. 5. Connectivity (probability of communication) and ignorance for steady-state simulations

However, it was not the aim of this section simply to suggest that the proximity protocol is necessarily a better choice for geosensor networks within the application domain of traffic networks. Instead, the preceding discussion has indicated how the simulation testbed can be used to begin to identify key strategies and factors affecting MAGNET performance. In turn, we expect that such investigations can help researchers and application domain experts begin to understand the behavior of dynamic geosensor networks. The preliminary results are a first attempt at evaluating different communication strategies, but in order to derive general recommendations further work is needed to ensure the results reflect more realistic environments and agent behavior. The following list summarizes a few related results of further simulations in this particular application domain:

- The proximity communication strategy is, as expected, sensitive to the threshold d , the distance beyond which agents using the proximity strategy will not inform other agents about a hazard. Further investigations indicated that d can be relatively small, in comparison with the environment, and still achieve effective communication.
- Simulations where agents employ a mix of different strategies often achieve high levels of efficiency. In particular, a small proportion of agents employing the flooding communication strategy mixed into a majority of agents using proximity or epidemic strategies can achieve low levels of agent ignorance and redundancy.
- Using a regular grid as a simulation environment is clearly an oversimplification with respect to real transportation networks. Simulations using environments without such uniform transportation network yield different results.

Hazards placed in hard-to-reach regions of the transportation network generally result in higher levels of ignorance and require longer to equilibrate. Conversely, hazards placed in easy-to-reach high traffic-density regions of the transportation network stabilized rapidly with almost perfect information dispersal (extremely low levels of ignorance).

6 Discussion and Conclusions

In this paper, we presented a simulation environment for testing information dissemination strategies in MAGNETs. We proposed and evaluated several strategies for scalable peer-to-peer information exchange (i.e., flooding-based, epidemic, and location-constrained). Strategy performance was measured based on the level of ignorance, redundancy, and degree of redundancy. Our simulation results showed that the proximity communication strategy provides an efficient compromise in terms of information dissemination in MAGNET efficiency. The levels of agent ignorance achieved using the proximity strategy are comparable to those achieved using the flooding strategy. At the same time, the proximity strategy does not lead to as high levels of information redundancy as the flooding strategy.

Several important issues have not been addressed in this paper, and will need to be the subject of further research. We identify below three core areas of related future research:

- Issues of *privacy* have not yet been addressed. The spatiotemporal reference associated with sensed information could be used to infer the location and movements of the agent(s) that sensed that information. Where the agents in our MAGNET are people, the location of those people becomes personal information that should not be widely disseminated. The development of strategies for safeguarding personal privacy in a decentralized geosensor network is a core issue for future research in this area.
- In this research we have assumed that the information generated by a geosensor is precise and accurate. In reality, *uncertainty* is an endemic feature of spatial and spatiotemporal information. The situation is further complicated by the possibility of malicious agents deliberately spreading misinformation throughout the geosensor network. To be practical, MAGNETs must be robust enough to continue to operate in the face of uncertainty arising from whatever source. In particular, the ability to resolve inconsistencies between multiple contradictory items of information from different sources is vital.
- While the decentralized model offers many advantages for mobile and location-aware systems, it may not be suitable for all types of location-based services. For example, safety critical applications may only be able to tolerate zero or minimal levels of ignorance. A decentralized ambulance routing location-based service, for example, could not operate if ambulances received the information they required only 90% of the time. In such an application, the minimal ignorance afforded by a centralized architecture would be vital,

possibly even at the cost of decreased system scalability and performance. Future research will need to address the suitability of the MAGNET architecture to specific application domains, perhaps developing hybrid approaches, where critical information is disseminated using a centralized server, while non-critical information is disseminated using a decentralized model.

Acknowledgments

This work was partially supported by the National Science Foundation under NSF grant number EPS-9983432 and the National Geospatial-Intelligence Agency under NGA grant number NMA201-01-1-2003. The authors would like to thank the anonymous reviewers for their helpful and constructive comments.

References

1. Kurose, J.F., Ross, K.W.: Computer Networking. 3rd edn. Addison Wesley, Reading, MA (2004)
2. Mennecke, B.E., Strader, T.J.: Mobile Commerce: Technology, Theory and Applications. Idea Group Publishing (2003)
3. Nittel, S., Stefanidis, A., Cruz, I., Egenhofer, M., Goldin, D., Howard, A., Labrinidis, A., Madden, S., Voisard, A., Worboys, M.: Report from the first workshop on geo sensor networks. Volume 33. (2004) 141–144
4. Hightower, J., Borriello, G.: Location systems for ubiquitous computing. *Computer* **34** (2001) 57–66
5. Weiser, M.: The computer of the 21st century. Volume 265. (1991) 94–104
6. Committee on Networked Systems of Embedded Computers, National Research Council: Embedded, Everywhere: A Research Agenda for Networked Systems of Embedded Computers. National Academies Press, Washington, DC (2001)
7. Crossbow Technology Inc.: MICA2DOT specification. <http://www.xbow.com> (2004) Last accessed: 03 July 2004.
8. Kahn, J., Katz, R., Pister, K.: Next century challenges: Mobile networking for smart dust. In: Proceedings of the 5th International Conference on Mobile Computing and Networking (MOBICOM), Seattle, Washington, USA. (1999) 271–278
9. Priyantha, N.B., Chakraborty, A., Balakrishnan, H.: The cricket location-support system. In: Proceedings of the 6th Annual ACM International Conference on Mobile Computing and Networking (MOBICOM00), Boston, MA. (2000) 32–43
10. University of California Los Angeles: WINS: Wireless integrated network sensors. <http://www.janet.ucla.edu/WINS> (2004)
11. Mainwaring, A., Polastre, J., Szewczyk, R., Culler, D., Anderson, J.: Wireless sensor networks for habitat monitoring. In: Proceedings of the 1st International Workshop on Wireless Sensor Networks and Applications (WSNA02), Atlanta, GA. (2002) 88–97
12. Ailamaki, A., Faloutsos, C., Fischbeck, P.S., Small, M.J., VanBriesen, J.: An environmental sensor network to determine drinking water quality and security. *ACM SIGMOD Record* **32** (2003) 47–52
13. Imielinski, T., Goel, S.: DataSpace: Querying and monitoring deeply networked collections in physical space. In: Proceedings of the 1st ACM International Workshop on Data Engineering for Wireless and Mobile Access (MobiDE'99), Seattle, Washington, USA. (1999) 44–51

14. Intanagonwiwat, C., Govindan, R., Estrin, D., Heidemann, J., Silva, F.: Directed diffusion for wireless sensor networks. *ACM/IEEE Transactions on Networking* **11** (2002) 2–16
15. Braginsky, D., Estrin, D.: Rumor routing algorithm for sensor networks. In: *Proceedings of the 1st International Workshop on Wireless Sensor Networks and Applications (WSNA02)*, Atlanta, GA. (2002) 22–31
16. Madden, S., Franklin, M., Hellerstein, J., Hong, W.: TAG: Tiny AGgregate queries in ad-hoc sensor networks. In: *Proceedings of the USENIX Symposium on Operating Systems Design and Implementation*, Boston, MA. (2002) 131–146
17. Estrin, D., Govindan, R., Heidemann, J.: Next century challenges: Scalable coordination in sensor networks. In: *Proceedings of the 5th International Conference on Mobile Computing and Networking (MOBICOM99)*, Seattle, Washington, USA. (1999) 263–270
18. Krishnamachari, B., Estrin, D., Wicker, S.: Modelling data-centric routing in wireless sensor networks. *Computer Engineering Technical Report CENG 02-14*, University of Southern California (2002)
19. Schiller, J., Voisard, A.: Information handling in mobile applications: A look beyond classical approaches. In Stefanidis, A., Nittel, S., eds.: *GeoSensor Networks*. Taylor and Francis (2004) Forthcoming.
20. Wolfson, O., Ouksel, A., Xu, B.: Resource discovery in disconnected mobile ad-hoc networks. In: *Proceedings of the International Workshop on Next Generation Geospatial Information*, Boston, MA. (2003)
21. Xu, B., Ouksel, A., Wolfson, O.: Opportunistic resource exchange in inter-vehicle ad-hoc networks. In: *Proceedings of the IEEE International Conference on Mobile Data Management (MDM'04)*, Berkeley, CA, USA. (2004) 4–12
22. Wolfson, O., Xu, B., Sistla, A.P.: An economic model for resource exchange in mobile peer to peer networks. In: *Proceedings of the 16th International Conference on Scientific and Statistical Database Management (SSDBM'04)*, Santorini, Greece. (2004) 235–244
23. Broch, J. and Maltz, D.A., Johnson, D.B., Hu, Y.C., Jetcheva, J.: A performance comparison of multi-hop wireless ad hoc network routing protocols. In: *Proceedings of the 4th annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM)*, Dallas, Texas, USA. (1998) 85–97
24. Johansson, P., Larsson, T., Hedman, N., Mielczarek, B., Degermark, M.: Scenario-based performance analysis of routing protocols for mobile ad-hoc networks. In: *Proceedings of the 5th annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM)*. (1999) 195–206
25. Raju, J., Garcia-Luna-Aceves, J.: A comparison of on-demand and table driven routing for ad-hoc wireless networks. In: *Proceedings of the IEEE International Conference on Communications (ICC 2000)*. Volume 3. (2000) 1702–1706
26. Der Spiegel: Sprechende Autos: Zu schlau fuer den Stau. <http://www.spiegel.de/auto/werkstatt/0,1518,267313,00.html> (2003) Last accessed: 03 July 2004.
27. Franklin, S., Graesser, A.: Is it an agent or just a program? *Proceedings of the 3rd International Workshop in Agent Theories, Architectures and Languages* (1996) 21–35
28. Sperber, D., Wilson, D.: *Relevance: Communication and Cognition*. Blackwell Publishers, 2nd Edition (1995)