Sensor Data Dissemination through Ad Hoc Battlefield Communications*

Linda Briesemeister
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025, U.S.A.
linda.briesemeister@sri.com

Keywords Mobile ad hoc network, sensor network, mobility model, battlefield scenario

Abstract

We study the dissemination of sensor data (reports) from the sensor network to the mobile ground forces (soldiers) for sensor gateways deployed in a battlefield scenario. Our approach looks at both the addressing and distribution of sensor reports.

First, we employ a subscription mechanism in which the sensor gateways address their reports to those soldiers who have currently subscribed to them. Second, we distinguish two schemes for propagating sensor reports. In a centralized approach, all sensor reports must go through one designated node (command post). In a distributed approach, the network routes sensor reports directly to the soldiers.

In a generic soldier mobility model, soldiers move in small groups (squads) along a line to random destinations on the battlefield. Through simulations using this mobility model, we study the performance and overhead of the proposed methods for sensor data dissemination.

We envision this research to be the first of a series of methods to manage information within mobile networks comprised of sensors and actuators in battlefield scenarios.

INTRODUCTION

Different types of wireless networks can occur in a battlefield scenario. Sensor networks contain many nodes often with small transmission ranges and limited battery lives. Unless deployed on vehicles or with ground forces, the sensors remain immobile. When equipped with wireless communications enabling them to exchange messages directly, vehicles and ground forces form mobile ad hoc networks that operate without an existing network infrastructure. In mobile ad hoc networks, the communication devices typically cover larger ranges than in sensor networks. Also, energy efficiency—albeit of concern—is not as critical for mobile ad hoc networks as for small sensor nodes. However, the definitions of these two network types overlap and hence a clear distinction cannot always be made.

While sensor networks themselves draw many research activities [1, 2, 3] we focus in this work on the higher level of disseminating the information gathered by the sensor network. We assume sensors are in fixed locations. So-called sensor gateways collect the raw data from a sensor network and provide an interface to the ad hoc network of mobile forces in the battlefield. The gateways may have post-processed the data before they send out sensor reports. In our model, the gateways are also immobile.

Many aspects of the interaction between sensor and mobile ad hoc networks deserve interest from a research perspective. An adjudication process gathers votes about the interpretation of data. Aggregation of data leads to a more complete view of the situation. Outside knowledge can be injected to augment automatic classification of sensor data. As a first step toward a complex model for handling sensor data in a mobile network, we focus on dissemination strategies of sensor reports from gateways to mobile forces.

In this article, we integrate the proposed dissemination strategies with a protocol stack of nodes in mobile ad hoc networks. We combine this system with the network simulation and a realistic model for pedestrian group mobility.

DISSEMINATION AND ADDRESSING METHODS

Our system delivers sensor reports to mobile forces in the field. Generally, soldiers only want to receive sensor reports relevant to their interest. Reducing the number of recipients to only interested entities can lead to less data traffic, which in turn can improve the throughput and end-to-end delay of data transmissions. In our model, the proximity of soldiers to

*The work reported in this paper was supported by DARPA under Contract No. F30602-02-C-0135.

Table 1. Combination of different methods

		Addressing	
		“Selective”	“Flooding”
Dissemination	Centralized	Soldiers subscribe at command post. Sensor gateways send reports to command post. Command post distributes reports to all subscribers.	Sensor gateways send reports to command post. Command post sends reports to all soldiers. Soldier’s devices decide upon receiving report, if relevant.
	Distributed	Soldiers subscribe at sensor gateways. Sensor gateways send reports to all subscribers.	Sensor gateways send reports to all soldiers. Soldier’s devices decide upon receiving report, if relevant.

sensor fields determines their interest in acquiring that data.

We assume fixed sensors and gateways. Soldiers know their own location via the Global Positioning System (GPS) and obtain a list of sensor gateways and their locations prior to entering the battlefield.

We employ a simple subscription mechanism to allow the addressing of sensor reports. Each soldier subscribes to his or her closest sensor gateway. The soldier’s device monitors the position constantly and compares the distances to all sensor gateways. When the affiliation with a gateway changes, the soldier sends two messages: One to unsubscribe from the old gateway and one message to subscribe to the new gateway. We call this approach “selective” addressing.

We compare the subscription method with an approach using pure flooding. There, we omit the subscription procedure. Instead, the sensor reports flood the network and soldiers receiving a copy compare the origin of the report with their current interest (= closest sensor gateway). If the received report came from the closest gateway, the device delivers the report to the application for further processing. If the sensor report originated from outside the soldier’s current interest, the device discards the report. We also studied this type of anonymous addressing through flooding in the context of highly mobile ad hoc networks [4, 5].

For propagation of sensor reports, we compare two schemes. In the centralized approach, all sensor reports must go through a central authority (command post). Hence, gateways send their reports to the command post and the command post forwards them to the soldiers (either to subscribers or through flooding). The soldiers in this scheme send their subscription requests to the command post, which keeps track of them. In the distributed approach, the sensor reports go directly to the soldiers. Consequently, the soldiers also send their subscription messages directly to the respective gateways.

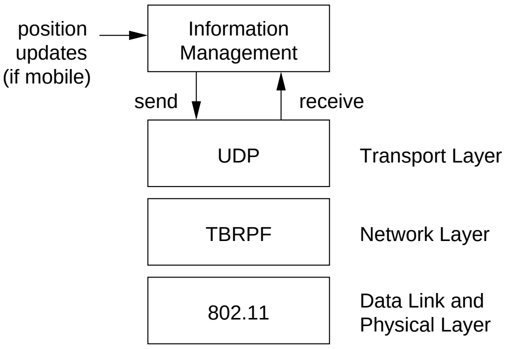

Figure 1. Layered system architecture

Table 1 summarizes the combinations of addressing and dissemination methods.

Simulation Model

We distinguish three types of nodes: Sensor Gateways, Soldiers, and the Command Post. In the simulation model, the top layer implements the behavior of these node types. Figure 1 shows the layered architecture of the system model.

The information management layer uses User Datagram Protocol (UDP) transmissions for sensor reports and subscription messages. The network layer resides below the transport layer. We choose our Topology Broadcast Based on Reverse-Path Forwarding (TBRPF) protocol [6, 7] for routing because it suits mobile ad hoc networks of medium size well.

We carry out the simulations with the ns-2 simulator [8, 9]. For wireless networks, ns-2 offers a model of the IEEE 802.11 medium access control and physical layer.

Unfortunately, our simulation implementation lacks support for multicasting. When dispatching the same sensor report to multiple soldiers, one would favor a native multicast transmission. Instead, we send the report via plain unicast to multiple soldiers. A metric then captures a theoretical benefit of a multicast by looking at the resulting paths of the individual unicasts corresponding to one report.

Sensor Gateway Distribution

We distribute the sensor gateways randomly over the battlefield. We distinguish two configurations with 20 and 40 gateways to resemble a sparse and dense network respectively.

We accept sparse configurations (with 20 gateways) only if the network topology of the sensor gateways with the fixed command post breaks at least once leaving the network fragmented. Conversely, we accept dense configurations (with 40 gateways) only if the network topology of sensor gateways is connected so that a path from every node to every other node exists.

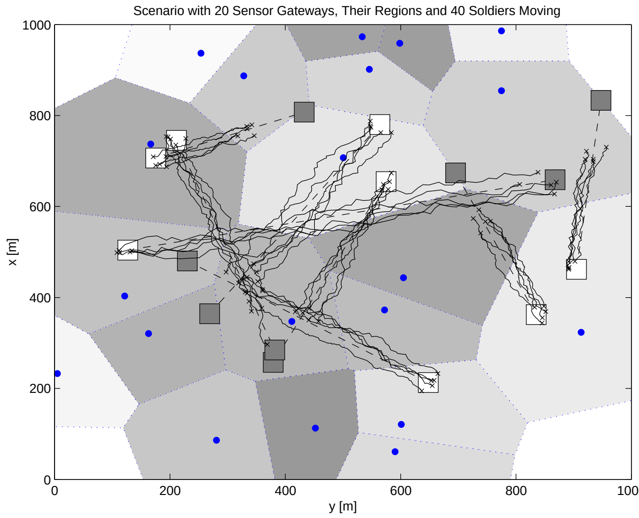

Figure 2. Example scenario for sensor gateway distribution and soldier mobility

Soldier Mobility

We apply a generic but fairly realistic soldier mobility model to our simulation scenarios. In this model, soldiers move in small groups of 5 (squads) along imaginary lines to random destinations on the battlefield. The soldiers' velocity stays at or below pedestrian speed of 2 meters per second. We draw the initial positions of a squad randomly from a 40 m-by-40 m square around the starting point of the imaginary line. Then, every 10 s, each soldier chooses a random angle within $\frac{\pi}{4}$ of the current direction towards his or her ultimate destination. We look at scenarios with 20 and 40 soldiers.

Figure 2 shows an example of a scenario with 20 sensor gateways and 40 soldiers. The soldiers are moving in 8 squads of 5 people each along a line (dashed) between random start (white big squares) and ultimate end (grey big squares) positions of their squad.

Parameters

Other parameters of the simulation model are as follows. The battlefield has a size of 1000 m-by-1000 m square with the command post placed close to the top left corner at (250,750). We simulate 600 s of real time. The gateways generate sensor reports every 5 seconds. For the IEEE 802.11 communication devices, we choose the parameters suggested in the setup files for ns-2: 2 Mb/s bandwidth, energy levels that translate to 250 m transmission range.

We produced five different scenarios for each sensor gateway distribution and each soldier mobility model. Combining the settings of 20 and 40 gateways with 20 and 40 soldiers that results in $5 \times 5 \times 2 \times 2 = 100$ configurations. We augmented every configuration with five different seeds for the random number generator. Finally, we ran each configuration with the centralized and distributed approach for sensor data dissemination obtaining $100 \times 5 \times 2 = 1000$ runs in total.

Table 2. Results for UDP transmissions

Gateways	Soldiers	Centralized		Distributed	
		Throughput	Delay	Throughput	Delay
20	20	37.48 kB/s	127 ms	84.17 kB/s	15 ms
20	40	28.42 kB/s	172 ms	73.63 kB/s	21 ms
40	20	35.50 kB/s	65 ms	89.49 kB/s	13 ms
40	40	28.29 kB/s	107 ms	83.50 kB/s	15 ms

Addressing sensor reports anonymously using the "flooding" approach, the simulation runs exceeded the computer resources. Therefore, we report herein only the results for 1000 runs with "selective" addressing.

Metrics and Results

Two types of UDP transmissions occur in selective addressing. The soldiers send subscription messages to either the command post (centralized) or directly to the gateway of interest (distributed). The sensor reports comprise the second type of UDP traffic. Table 2 gives the results for throughput and end-to-end delay of UDP transmissions.

The results for both metrics reflect that the transmission paths in the centralized approach are on average longer than in the distributed approach. The transmission range of a gateway covers most of the region associated with it. Then, the gateway reaches the interested soldiers using only one hop in the distributed approach. In contrast, messages through the command post take multiple hops reducing throughput and increasing delay.

Within each dissemination method, the throughput and delay metrics decline for configurations with more soldiers but stay similar for scenarios with more gateways. We explain this by using inefficient unicasts to propagate the same report to many recipients. The multiple unicasts—issued almost simultaneously—compete over network resources while the network treats them independently. Scenarios with more soldiers increase this effect of stress on the network. We expect a lesser difference for increasing the number of soldiers when applying native multicast.

We capture the performance of our dissemination methods using a metric of success. We distinguish a "requested success" from an "actual success" of one sensor report. For the requested success, we divide the number of of recipients of a report by the number of soldiers having *requested* subscription at the time the gateway issues the report. For actual success, we divide the number of recipients of a report by

Table 3. Results for sensor reports

Gateways	Soldiers	Centralized		Distributed	
		Requested success	Actual success	Requested success	Actual success
20	20	78.51 %	93.99 %	99.16 %	99.72 %
20	40	85.89 %	93.20 %	99.21 %	99.84 %
40	20	99.53 %	99.14 %	99.94 %	99.99 %
40	40	99.38 %	99.17 %	99.93 %	99.99 %
		Control overhead	Unicast overhead	Control overhead	Unicast overhead
20	20	2.98 %	55.33 %	5.79 %	64.93 %
20	40	3.39 %	61.52 %	4.85 %	70.34 %
40	20	2.45 %	52.31 %	7.77 %	59.07 %
40	40	3.23 %	55.74 %	6.66 %	63.39 %

the number of subscribers that the command post or gateway knew at the time of it issued the report.

Table 3 summarizes the results for performance and overhead metrics of the dissemination methods.

The success for the centralized approach is always less than for the distributed approach. We explain this by looking at the paths messages take. In the centralized approach, subscription and sensor messages go through the command post taking multiple hops. Recall that for configurations with 20 gateways we forced the fixed topology to be disconnected. Hence, the network cannot always propagate subscription reports and sensor reports to the desired destination. The requested success suffers especially from counting those soldiers who sent a subscription message that never reaches the command post so that the list of subscribers is possibly incomplete.

Not surprisingly, the fragmentation of the network does not hamper the performance of the distributed method. Here, most of the subscription messages and sensor reports take only one hop to their destination making the transmission less likely to fail.

The next metric considers the overhead caused by applying the subscription method. The so-called control overhead is the percentage of subscription packets[1] of the overall amount of packets.

The control overhead stays below 7.77 % for all configurations. Note that the mobile nodes move at pedestrian speed. Therefore, the event of a soldier crossing the boundary between sensor regions occurs less often than in a scenario with faster ground vehicles. We also plan to include overlapping regions in our future studies, which will introduce more subscription traffic.

The control overhead of the centralized approach is smaller than for the distributed approach. But the amount of non-subscription packets in the centralized approach is significantly greater than in the distributed approach, because all sensor reports go through the command post taking more hops than sensor reports in the distributed approach. Therefore, the ratio of subscription packets to all packets causes a smaller control overhead for the centralized approach.

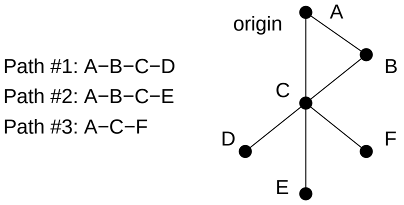

Figure 3. Example of calculating unicast overhead

The last metric assesses the impact of using plain unicast instead of native multicast for sending the same report to multiple recipients. For the denominator, we count the number of hops of all unicast messages belonging to one sensor report. Then, we parse all paths in a linear fashion counting only those edges that none of the paths already parsed included. Also, for the last hop before delivering the report to the recipient, we only count this edge if the outgoing node is new. Thus, we also capture the benefit of radio communication with an omnidirectional antenna. Finally, we obtain the unicast overhead by subtracting the number of unique edges (without the last hop from known nodes) from the denominator and then dividing by the denominator.

Figure 3 illustrates an example of calculating unicast overhead. Origin A sends three unicast messages to the destinations D, E, and F. The denominator here equals eight for the sum of all hops. Going through the paths one by one gives us the following numbers. Path #1 contains three new edges. Path #2 contains only one new edge C–E, but this is the last hop and the edge comes from a known node C. Therefore, we add zero to the number of unique edges. Path #3 yields one new edge A–C. Again, the last hop C–F originated from a known node so that Path #3 adds one to the number of unique edges, The unicast overhead for this example then is $8 - (3 + 0 + 1)/8 = 0.5$.

The unicast overhead attempts to estimate the savings of using multicast over multiple unicasts. However, an implementation of multicast can result in more or less savings then this metric estimates.

The unicast overhead of the distributed method is always greater than for the centralized approach. In the distributed approach, the unicast messages take, in most cases, one hop to the destination. Ideally, only one transmission is then necessary to reach all destinations in range of an omnidirectional antenna. In the centralized approach, the routing layer sometimes chooses multihop paths that have minor differences. But our linear parsing of known edges may not account for a

[1] One packet equals one hop on a path of a message transmission.

slightly different route that uses new edges. Like the example above, taking a shortcut in Path #3 does not revert counting the edges A–B and B–C in Paths #1 and #2 and instead decreases the difference between numerator and denominator in our formula. Then, the unicast overhead is smaller than for one hop multicasts.

Not surprisingly, the unicast overhead increases with the number of soldiers because they represent the destinations of sensor reports being split up into multiple unicasts.

CONCLUSION

We describe an ad hoc network protocol stack for disseminating sensor data in a battlefield scenario. Simulations of the system with a realistic mobility model of soldier movement let us compare a centralized with a distributed approach.

The simulation results show a better performance of the distributed approach. The distributed method benefits from the locality of the task to deliver reports to soldiers in proximity. Network disconnections have less impact on reaching the subscribers of sensor reports nearby. Compared to the distributed method, the centralized approach causes slightly less control and less unicast overhead. However, the control overhead measures the relation of control traffic to all traffic, which is significantly larger using the centralized method. An implementation of multicast will mend the problem of unicast overhead.

Still, the centralized approach bears benefits that are not accounted for in our metrics, but may be important in future information management. The central authority of the command post can provide adjucation and decision power. It also allows easy access to the network to inject outside knowledge. In addition, it reflects existing command hierarchies in military applications. To overcome the problem of network fragmentation, long-haul links can provide a fall back level for communication devices if the network suffers from disconnection to the command post.

For future research, we plan to implement multicast to better support such applications of ad hoc networks. We also want to improve the scenarios toward becoming more realistic. Here, sensor regions will overlap and the network will contain heterogenous node types ranging from ground vehicles to soldiers. We envision this research to be the first of a series of methods to manage information within mobile networks comprised of sensors and actuators in battlefield scenarios.

References

[1] Wendi Rabiner Heinzelman, Joanna Kulik, and Hari Balakrishnan. Adaptive protocols for information dissemination in wireless sensor networks. In *Mobile Computing and Networking*, pages 174–185, 1999.

[2] Alvin Lim. Distributed services for information dissemination in self-organizing sensor networks. *Special Issue on Distributed Sensor Networks for Real-Time Systems with Adaptive Reconfiguration. Journal of Franklin Institute*, 338(6):707–727, September 2001. ISSN 0016-0032.

[3] Sung Park, Andreas Savvides, and Mani B. Srivastava. Simulating networks of wireless sensors. In *2001 Winter Simulation Conference*, 2001.

[4] Linda Briesemeister, Lorenz Schäfers, and Günter Hommel. Disseminating messages among highly mobile hosts based on inter-vehicle communication. In *IEEE Intelligent Vehicles Symposium*, pages 522–527, October 2000.

[5] Linda Briesemeister. *Group Membership and Communication in Highly Mobile Ad Hoc Networks*. PhD thesis, School of Electrical Engineering and Computer Science, Technical University of Berlin, Germany, November 2001.

[6] Bhargav Bellur and Richard G. Ogier. A reliable, efficient topology broadcast protocol for dynamic networks. In *IEEE Infocomm*, pages 178–186. IEEE, March 1999.

[7] Richard G. Ogier, Fred L. Templin, Bhargav Bellur, and Mark G. Lewis. Topology broadcast based on reverse-path forwarding, November 2002. IETF Internet Draft (work in progress). http://www.ietf.org/internet-drafts/draft-ietf-manet-tbrpf-06.txt.

[8] Steven McCanne and Sally Floyd. ns Network Simulator. http://www.isi.edu/nsnam/ns/.

[9] Lee Breslau et al. Advances in network simulation. *IEEE Computer*, 33(5):59–67, May 2000.

Linda Briesemeister received her graduate degree in Computer Science in March 1998 and Ph.D. degree in Engineering in November 2001, both from the Technical University of Berlin, Germany. She is a Research Engineer in the Information, Telecommunications, and Automation Division (ITAD) of SRI International. Her research includes distributed algorithms applied to mobile computing and wireless communication as found in ad-hoc networking. Prior to joining SRI, she worked on intervehicle communication at DaimlerChrysler Research and Technology. Dr. Linda Briesemeister is a member of the ACM.