Efficient Group Mobility for Heterogeneous Sensor Networks

Patrick Traynor[1], JaeSheung Shin[1], Bharat Madan[2], Shashi Phoha[2], Thomas La Porta[1]

[1]Networking and Security Research Center
Department of Computer Science and Engineering
The Pennsylvania State University

[2]Information Science and Technology
Applied Research Laboratory
The Pennsylvania State University

***Abstract*— Mobility management protocols allow wireless devices to move between networks. These protocols have traditionally supported the mobility of individual nodes and are therefore not optimized to support the migration of groups. Accordingly, the time required to re-establish connectivity, frequency of dropped packets and contention for the air interface increase significantly for mobile groups. We propose a protocol for mobile groups that reduces all of the above by allowing a single node to perform handoffs on behalf of all group members. This "gateway" node eliminates the need for multiple handoff messages by obscuring group membership to external parties. Through extensive simulation and implementation, we show significant reduction in handoff times, message complexity and packet loss for groups of heterogeneous, mobile sensors running AODV and DSDV. By leveraging the naturally occurring hierarchy, we demonstrate that it is possible for groups to *efficiently* use traditional mobility protocols to support their collective movements.**

Keywords: Group Mobility, Hierarchical Routing, Heterogeneous Sensor Networks, Network Address Translation

I. Introduction

In order to provide required coverage or track moving units, many proposed sensor networks incorporate mobility. As a node moves across a number of access points or data sinks, it performs a number of handshaking exchanges that allow traffic generated by and for it to be delivered nearly seamlessly. The movement of nodes, however, is not always independent. A convoy of trucks on the highway, a sightseeing bus full of tourists and a platoon of marching soldiers all represent realistic examples of mobile groups. Accordingly, the nearly simultaneous attempts of a large number of nodes negotiating service with a new access point have direct effects on connectivity and performance.

If a large number of related nodes attempt to perform the handshaking associated with a handoff within a short period of time, a number of problems arise. First, the number of collisions over the wireless interface is likely to be high, causing additional power drain to occur from supplementary transmissions. These collisions will also result in delay, which may cause problems for applications that are delay sensitive. There is little need in sending multiple handshaking messages for mobile node to a new access point if we expect the arrival of a group. If one node could represent the remainder of the group and take care of the handshaking messages with a new base station, a much more efficient mobility model would be created. We refer to this solution as *group mobility*.

We leverage naturally occurring hierarchies in heterogeneous mobile sensor networks to select nodes as "gateways" between their group and IP access points. These gateway nodes perform a Mobile IP handoff on behalf of all group members at each new point of attachment. These gateways also obscure the membership and topology of their groups to external entities by acting as network address translation (NAT) [7] devices. In so doing, this method allows groups of mobile nodes to use native protocols to manage communication and mobility within their group. This property makes the group mobility solution suitable for different types of mobile network, including ad hoc networks and networks of mobile sensors.

In this paper, we make the following contributions:

- **Protocol Definition:** We define the protocol necessary to perform efficient mobility management for a group of mobile nodes. While similar ideas are currently being examined by the IETF Network Mobilty (NeMo) Working Group [12], we are the first to implement, measure and characterize this idea for sensor networks.
- **Performance Analysis:** Using NS2 simulations and an implementation on a variety of platforms, we demonstrate the performance gains and practicality of group mobility.
- **Design Tradeoffs:** By tuning a number of options including the choice of ad hoc routing protocol and gateway selection, we demonstrate that the group mobility approach can be customized to the needs of specific networks.

The remainder of this paper is organized as follows: In Section II, we examine related work; in Section III, we define the mechanisms used to implement group mobility; Section IV explores the results of NS2 simulations and a mote-based implementation; Section V explores the pertinent issues for networks implementing this solution; Section VI offers concluding remarks and future work.

II. Related Work

Sensor networks have traditionally been modeled as static systems composed of hundreds of homogeneous, resource constrained devices [2]. Recent work in this area has challenged these assumptions. Networks leveraging mobility to compensate for the ad hoc nature of node deployment and the resulting poor coverage have been increasingly discussed [18], [19]. Other networks rely upon patrolling nodes to relay data from otherwise disconnected regions [20]. A number of researchers are also taking advantage of networks composed

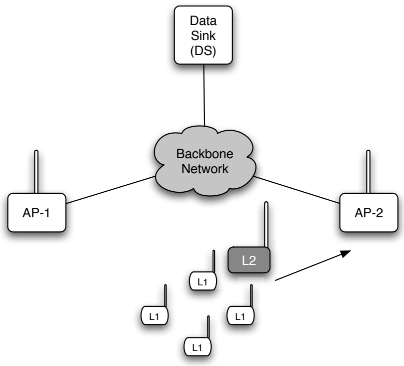

Fig. 1. An example of a heterogeneous group moving between access points (APs). As the signal strength of AP-2 becomes greater than that of AP-1, the L2 node performs a Mobile IP handoff, allowing all traffic for the group to be routed to its new location.

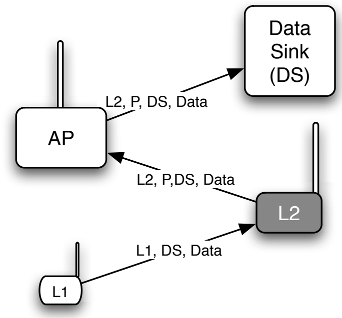

Fig. 2. A message is sent from an L1, translated by a NAT-like process in an L2 and forwarded from port P via an AP to its intended destination. Commands or requests from the data sink follow the reverse path, except that the traversal of the wired network may be different due to Mobile IP.

of heterogeneous platforms. Mhatre, et al. [11] designate nodes with access to additional resources as cluster heads so as to maximize the lifetime of their networks. Traynor, et al. [17] attempt to create more secure systems by implementing probabilistic keying and creating an unbalanced key distribution based on node capability. Through all of the above considerations, researchers are designing networks that are more able to perform in highly dynamic environments.

The inclusion of mobility and heterogeneity allow many of the protocols used in all IP networks to be incorporated into sensor networks. For example, as these systems grow in size, the mobility of the most capable nodes can be managed through Mobile IP [15], [10]. Packets destined for a targeted node are delivered to that node's "home" network, where they are forwarded on the current "foreign" location of that node based on a "care-of" address. This protocol, however, is designed to account for the mobility of an individual node and is therefore not optimized for groups. Because mobile groups and clustered movement are common in realistic settings [3], [4], [16], the mixed composition of these groups can used to create more robust systems.

Heterogeneity naturally breeds hierarchy in networks. Pei, et al. [14] attempt to take advantage of the available strata by breaking mobile groups into local subnets. Packets destined for subnet members are then routed through the hierarchy. Later work creates a more extensive model of the above in which nodes keep accurate depictions of group topology and route packets toward a destination's most significant "landmark" neighbors [13]. Similar techniques are currently being examined by the IETF Network Mobility (Nemo) Working Group [12]. By exploiting naturally occurring hierarchies, we can reduce the overhead incurred by traditional mobility protocols when used for groups.

By obscuring local topology, the performance of group handoffs is accordingly made equivalent to that of a single node. The details of the mechanisms necessary to accomplish this are discussed in the next section.

III. System Model

A. Network Components

Heterogeneity is common in almost every large system. It is this characteristic that makes populations both biological and mechanical in nature more robust to environmental changes. Accordingly, we deviate from the traditional model of sensor networks by assuming heterogeneous composition. In terms of wireless nodes, we create a simple two-level hierarchy. Depending on the needs of an application, this hierarchy can be extended as necessary.

The most resource constrained nodes in our model are the sensing, or Level 1 (L1) nodes. These nodes include such platforms as the Crossbow MICAZ [5]. Because of their simple architecture, these systems are unable to run many of the protocols used in IP-based networks. L1 nodes are the standard platforms used to create homogeneous sensor networks. Level 2 (L2) nodes represent platforms with significantly greater processing and storage capabilities. Because of their increased resources, these nodes serve as routers between IP networks and sensor (L1) platforms.

For the sake of simplicity, we consider the wired portion of this system to be composed of access points (APs), "home" and "foreign" agents and some backbone network responsible for moving packets to their intended destinations. Figure 1, above, gives a sample network containing these components.

B. Protocol Definition

For this particular system, we consider a network in which sensor nodes report data to some data sink in the backbone network. This does not preclude a data-centric approach to the wireless portion of this network.

Figure 2 illustrates message delivery. L1 nodes transmit sensed data at regular intervals. Because these nodes have extremely limited capabilities, packet delivery uses a lightweight, non-IP-based mesh networking stack. The data from each mote is received by an L2 and is placed in an IP packet. The L1 source and wired destination of the data are stored in a NAT [7]-like table so that any future responses from the wired network can be properly routed. Figure 3 shows this transformation. The L2 gateway then sends the packet to its serving AP. Data is forwarded through the backbone network to its intended location. Responses from the data sink follow the reverse of the route described here.

By using L2 nodes as NAT devices, packets can be delivered to all of the L1 nodes without requiring the wired portion of network to keep track of the specific topology of any group.

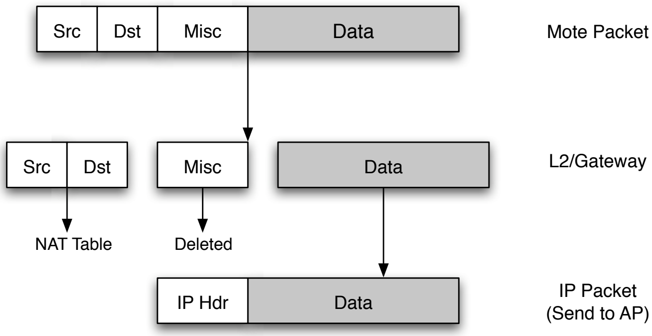

Fig. 3. An overview of the NAT-like packet processing in L2 nodes. Packets arrive at an L2 node on the interface corresponding to the mote networking protocol. The necessary data is removed and the process creates or checks a table entry corresponding to the source and destination addresses. The data is then placed in an IP packet and forwarded to its destination.

Likewise, the L1 nodes may run any native protocol to manage their local communications and mobility. For example, an ad hoc network of relatively capable devices can run an ad hoc routing protocol such as AODV. A sensor network comprised of very limited nodes may run protocols optimized for their platform and application.

The message complexity of this approach for managing mobility is greatly reduced from an approach in which each node manages their mobility individually. In the individual case, the number of mobility messages will be on the order of the number of nodes, n, in a group, or $O(n)$. With group mobility, the number of messages for managing mobility for handoffs between access points will be constant, independent of the number of nodes in a group, or $O(c)$.

IV. Experimental Results

In order to demonstrate that group mobility provides tangible improvements to mobile groups of sensors, we created simulations using NS2 and implemented a test system.

A. Simulation

The effects of group mobility on mobile sensor networks were examined using the NS2 simulator version 2.27. The effects of mobility were examined using the included DSDV and AODV-UU, a standards compliant version of AODV [1]. All nodes used an 802.11 MAC layer. Packets were sent every 0.05 seconds both from wired and wireless nodes. Nodes had a transmission radius of 30 meters; however, all group members were not necessarily within transmission range of each other. The transmission radii of APs overlap so that nodes are never in a region without service. Each simulation was run 50 times with a randomly generated topology for a single group. Group membership remains constant throughout these experiments.

We examine node movement according to two general classifications. The first, referred to as *formation*, represents nodes moving in a rigid pattern so as to preserve relative positioning between neighbors. This correspondis to a group of vehicles driving in formation down a road. The second, which we refer to as *random*, represents intra-group movement following a random-waypoint model of mobility. The nodes themselves move in the same vector, but specific inter-node spacing is not preserved. This classification may be more applicable to a group of soldiers moving across a battlefield. Mobility management itself is classified as *individual* if each node is responsible for performing its own handoffs or *Group* if one node speaks on behalf of the entire group.

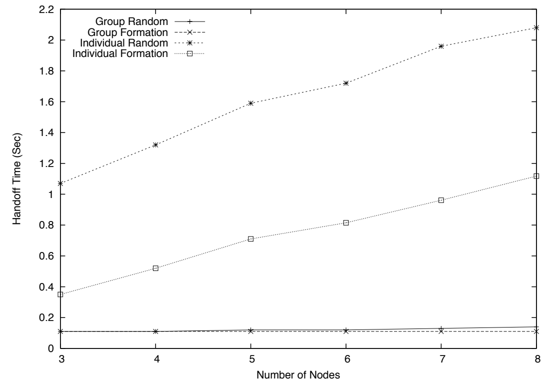

Fig. 4. The time required to perform successful handoffs for all nodes where the number of nodes per group varies. In these simulations, nodes were using DSDV to forward packets.

Figure 4 shows the average time required for each node in a group using DSDV to initiate a handoff and begin receiving packets from a new AP. Notice that both of the individual mobility times are much higher than the group mobility solution. As a group approaches a new AP, the propagation of connectivity with the backbone network takes an increasing amount of time as the number of nodes per group increases. Much of this overhead is due to the DSDV protocol itself. When using DSDV, all nodes are required to proactively update routing tables with next hop data to all other nodes in the network. Accordingly, Figure 4 shows the time required to begin receiving packets again for randomly moving nodes using individual mobility is approximately twice that of a cluster exhibiting rigid movement.

Figure 5 shows the results of similar tests on the reactive AODV protocol. Nodes performing handoffs individually are again significantly higher when compared to clusters using group mobility. However, the time required to perform a handoff is much lower than that of the system using DSDV. Because nodes only inquire about routes when they are needed, handoffs in AODV occur in the presence of much reduced background routing traffic. The use of the group approach still yields benefits as the path out of the network is always through a given node. Accordingly, nodes using reactive protocols like AODV must only keep track of the gateway node in order to remain connected with the wired network.

Figure 6 examines packet loss for both systems under the random mobility model. For both AODV and DSDV, the group mobility method reduced the number of dropped packets by an average of 5 and 10%, respectively. As expected from the previous results on handoff times, AODV performed better than DSDV. Packet loss for all scenarios became extremely high as the number of nodes per group increased past 15. This is due primarily to increased node density and a high sending rate.

B. Implementation

In order to better characterize the effects of group mobility, we have implemented a pulse monitoring application for

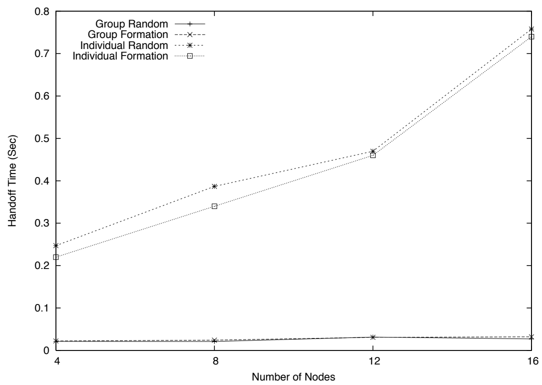

Fig. 5. The time required to perform successful handoffs for nodes using AODV to forward packets.

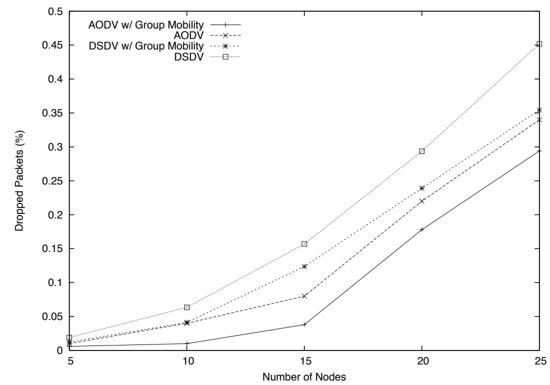

Fig. 6. The average packet loss for mobile nodes using the individual and group mobility mechanisms with the "random" movement model. For both AODV and DSDV, the group mobility method reduced dropped packets an average of 5 and 10%, respectively.

heterogeneous mobile platforms. As a group of soldiers patrols an area, each wears an L1 node that reports their current pulse to some central repository. One soldier, functioning as the "radio man", carries an L2 node. As the platoon patrols the perimeter of an area, the L2 node performs Mobile IP handoffs in order to keep the group's data flowing back to the central repository. In so doing, a central authority is able to keep track of the condition of all deployed soldiers without specifically requiring voiced radio contact.

Crossbow MICAz [5] motes, with a 4 MHz Atmel ATmega128L processor, 128 KB of program flash memory, and 512 KB of measurement flash memory are used as L1 nodes in this scenario. The TinyOS Oscilloscope program was used as the source of the generated "pulse" data. The L2 node is the combination of a MICAz mote mated with a Crossbow Stargate with a 400 MHz Intel PXA255 Xscale processor, 64 MB of SDRAM and 32 MB of flash memory. The Stargate node is equipped with an 802.11b wireless card and runs Dynamics Mobile IP [6] v0.8.1 for mobility management. The NAT functionality is implemented at the application layer, as the networking stacks of the L1 and L2 motes do not share common protocols at lower levels. All of the performance figures discussed below were collected from 5,000 iterations of the protocol.

Table I shows the results of microbenchmarks run on the Stargate L2 node. We specifically examine the additional overhead incurred by requiring data to be processed and stored in the NAT table. In this implementation, the packets arriving over the mote serial connection are parsed, the necessary data is placed into an IP packet, an entry in the NAT table is made or checked, and the packet is forwarded over the Stargate's 802.11 interface. Because this processing comprises less than 0.25% of the entire operation of the Stargate, the pre-processing associated with group mobility is not a bottleneck. Averaging 2.449 milliseconds (or 64.81% of the execution time), reading packets over the serial interface between the Stargate and its coupled mote will become the bottleneck as the number of nodes supported in a group grows.

V. Discussion

A. Additional Benefits

The previous sections of this paper have explored the group mobility protocol and many of its more obvious benefits. There are a number of additional benefits inherent to implementations of this protocol.

TABLE I
Microbenchmark results for group mobility implementation.

Operation	$\bar{x}$ Time (msec)	σ	%
Read Mote Packet	2.448	4.947	64.81%
Mote to IP Packet	0.005	0.0006	0.13%
NAT Lookup	0.004	0.0004	0.11%
Send IP Packet	1.320	3.649	34.95%

Mobility management in sensor networks is allowable chiefly because of the approach taken in this paper. The storage requirements to implement Mobile IP are far beyond the available 128KB of available program memory in a MICAz mote. Even if the Mobile IP codebase could be reduced in size, the radio stacks of available sensors are not IP-based. In order to use any of the "off-the-shelf" software suites of Mobile IP, an entirely new radio stack would have to be designed and implemented.

Heterogeneity of radio technology also makes this system more robust. Because L2 nodes are equipped with 802.11 cards, they are able to easily communicate at distances of 0.25 miles. With a larger, directed antennae, it is possible to increase this distance to tens of miles [8]. In contrast, the ZigBee equipped MICAz motes are limited to a transmission range of approximately 250 feet [5]. A platoon of soldiers carrying only L1 nodes would therefore have to remain much closer to APs as they made their patrols for connectivity to be maintained. Even if Mobile IP or equivalent code could be written for these constrained platforms, the hindering the movement of a group in this way would be unacceptable in a real system. The addition of the more powerful L2 node therefore makes the group as a whole more robust to environmental changes.

The group mobility solution allows nodes to talk to the nearest AP with low overhead. For individual mobility, in many a node may relay traffic through several hops to reach its original AP. Because the air interface is characteristically more lossy than wired connections, forcing packets to traverse a large number of hops wll increase the need for retransmissions. Even if the probability of packet loss at each node is calculated to be 5%, a 10-hop path would deliver just under 60% of the packets sent on it. By allowing nodes to seek shorter paths to closer APs, the frequency with which retransmissions occur can also be reduced.

B. Gateway Failure

Throughout this paper, we have assumed that the group mobility mechanism is executed in the presence of a single L2/gateway node. The difficulty with this approach is that it introduces a single point of failure in the system. If the L2 node were to be damaged or destroyed, communication by the remaining group members with the backbone network would likely be severed.

Real networks would therefore include multiple nodes capable of executing the duties of the gateway. The presence of multiple L2 nodes, however, necessitates additional overhead. L2 nodes must decide which among them will execute the functions of the gateway. If some subset of the L2 nodes act as gateways concurrently, the gains of the group mobility approach would be reduced. Regardless of whether one or more L2 nodes is currently acting as a gateway, their very presence forces L1 nodes to determine their corresponding gateway.

The simplest approach to solving this problem would be to have all of the L2 nodes function as a multicast group. These nodes would select an L2 to act as the serving gateway. While this node is active, the remaining nodes would ignore requests to forward data to the IP network. If the elected node should fail, a new L2 node could take the address of the former gateway (and therefore the address of the group) and allow service to continue. Because distributed election protocols are a widely studied area, we make the selection of specific algorithms as an orthogonal problem.

VI. Conclusion

Protocols including Mobile IP allow wireless devices to move between network access points while maintaining their IP addresses. While mobility protocols have traditionally focused on individuals, the movement of autonomous devices is often highly correlated. Recognizing this relationship allows the administration of such mobility to be made more efficient. In this paper, we discuss the application of the group mobility mechanism on heterogeneous, mobile sensor networks. By allowing a single node to perform Mobile IP registrations on behalf of an entire group, we demonstrate that it is possible to reduce messaging overhead for a mobile group from $O(n)$ to $O(c)$. In so doing, we have shown that we can decrease competition for the air interface, reduce the need for retransmissions and allow the flow of data to resume more quickly. Through the use of the available hierarchy, we demonstrate that groups can use unmodified versions of traditional mobility protocols to efficiently support their collective movements.

There are a number of directions for future work. First, the effects of data-centric protocols such as Directed-Diffusion [9] need to be analyzed. Additionally, we have not discussed the specific protocols used to create, join and leave mobile groups. These algorithms may have additional effects upon this protocol and should also therefore be examined.

VII. Acknowledgments and Disclaimer

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) and the National Science Foundation (NSF) and administered by the Army Research Office under ESP MURI Award No. DAAD19-01-1-0504 and the NSF WORKIT project (NSF Grant CNS-0453830). Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the Defense Advanced Research Projects Agency (DARPA), the Army Research Office or the National Science Foundation.

References

[1] Ad Hoc On Demand Distance Vector Routing - Uppsala University. http://core.it.uu.se/AdHoc/AodvUUImpl.

[2] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. A survey on sensor networks. *IEEE Communications Magazine*, August 2002.

[3] J. Cano and P. Manzoni. Group mobility impact over TCP and CBR traffic in Mobile Ad Hoc Networks. In *Proceedings of the IEEE Euromicro Conference on Parallel, Distributed and Network-Based Processing*, 2004.

[4] W. Chen and P. Chen. Group Mobility Management in Wireless Ad Hoc Networks. In *Proceedings of the IEEE Vehicular Technology Conference*, 2003.

[5] Crossbow Technologies. Wireless Sensor Networks. http://www.xbow.com/Products/Wireless_Sensor_Networks.htm.

[6] Dynamics Mobile IP. http://dynamics.sourceforge.net/.

[7] K. Egevang and P. Francis. RFC 1631 - The IP Network Address Translator. http://www.ietf.org/rfc/rfc1631.txt, 1994.

[8] P. Hernandez. Record Set at DefCon WiFi Shootout. http://www.enterpriseitplanet.com/networking/news/article.php/3524491, 2005.

[9] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed diffusion: A scalable and robust communication paradigm for sensor networks. In *Proceedings of the IEEE International Conference on Mobile Computing and Networking (MobiCOM)*, 2000.

[10] D. Johnson, C. Perkins, and J. Arkko. RFC 3775 - Mobility Support for IPv6. http://www.faqs.org/rfcs/rfc3775.html, 2004.

[11] V. Mhatre, C. Rosenberg, D. Kofman, R. Mazumdar, and N. Shroff. A Minimum Cost Heterogeneous Sensor Network with a Lifetime Constraint. *IEEE Transactions on Mobile Computing*, 2004.

[12] Network Mobility (NEMO) Working Group. http://www.ietf.org/html.charters/nemo-charter.html.

[13] G. Pei, M. Gerla, and X. Hong. LANMAR: Landmark Routing for Large Scale Wireless Ad Hoc Networks with Group Mobility. In *Proceedings of the ACM Symposium on Mobile ad hoc networking and computing (MobiHOC)*, 2000.

[14] G. Pei, M. Gerla, X. Hong, and C. Chiang. A Wireless Hierarchical Routing Protocol with Group Mobility (WHIRL). In *Proceedings of the IEEE Wireless Communications and Networking Conference*, 1999.

[15] C. Perkins. RFC 3220 - IP Mobility Support for IPv4. http://www.faqs.org/rfcs/rfc3344.html, 2002.

[16] G. Roman, Q. Huang, and A. Hazemi. Consistent Group Membership in Ad Hoc Networks. In *Proceedings of the IEEE International Conference on Software Engineering (ICSE)*, 2001.

[17] P. Traynor, H. Choi, G. Cao, S. Zhu, and T. L. Porta. Establishing Pair-Wise Keys in Heterogeneous Sensor Networks. In *Proceedings of IEEE INFOCOM*, 2006.

[18] G. Wang, T. L. Porta, and G. Cao. Movement-Assisted Sensor Deployment. In *Proceedings of IEEE INFOCOM*, 2004.

[19] G. Wang, T. L. Porta, G. Cao, and W. Zhang. Sensor Relocation in Mobile Sensor Networks. In *Proceedings of IEEE INFOCOM*, 2005.

[20] W. Zhao, M. Ammar, and E. Zegura. A Message Ferrying Approach for Data Delivery in Sparse Mobule Ad Hoc Networks. In *Proceedings of the ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHOC)*, 2004.