ASCENT: Adaptive Self-Configuring Sensor Networks Topologies

K.Seena Naik, Dr. Ramachandra, M.V.Brahmananda Reddy

I. INTRODUCTION

The availability of micro-sensors and low-power wireless communications will enable the deployment of densely distributed sensor/actuator networks for a wide range of environmental monitoring applications from urban to wilderness environments; indoors and outdoors; and encompassing a variety of data types including acoustic, image, and various chemical and physical properties [24]. The sensor nodes will perform significant signal processing, computation, and network self-configuration to achieve scalable, robust and long-lived networks [1, 7, 8]. More specifically, sensor nodes will do local processing to reduce communications, and consequently, energy costs.

These requirements pose interesting challenges for networking research. One of the challenges arises from the greatly increased level of *dynamics***.** The large number of nodes will introduce increased levels of *system dynamics*, which in combination with the high level of *environmental dynamics* will make designing reliable systems a daunting task. Perhaps the most important technical challenge arises from the *energy constraints* imposed by *unattended systems*. These systems must be long-lived and operate without manual intervention, which implies that the system itself must execute the measurement and adaptive configuration in an energy constrained fashion. Finally, there are *scaling* challenges associated with the large numbers of nodes that will co-exist in such networks to achieve desired spatial coverage and robustness.

In this paper, we describe and present experimental performance studies for a form of adaptive self-configuration designed for sensor networks. As we argue in Section 2, such unattended systems will need to self-configure and adapt to a wide variety of environmental dynamics and terrain conditions. These conditions produce regions with non-uniform communication density. We suggest that one of the ways system designers can address such challenging operating conditions is by deploying redundant nodes and designing the system algorithms to make use of that redundancy over time to extend the systems life. In ASCENT, each node assess its connectivity and adapts its participation in the multi-hop network topology based on the measured operating region. For instance, a node:

- Signals when it detects high message loss, requesting additional nodes in the region to join the network in order to rely messages to it.
- Reduces its duty cycle if it detects high message losses due to collisions.
- Probes the local communication environment and does not join the multi-hop routing infrastructure until it is "helpful" to do so.

Why can this adaptive configuration not be done from a central node? In addition to the scaling and robustness limitations of centralized solutions, a single node cannot directly sense the conditions of nodes distributed elsewhere in space. Consequently, other nodes would need to communicate detailed information about the state of their connectivity in order for the central node to determine who should join the multi-hop network. In the absence of energy constraints, one can always achieve a result that is closer to optimal with a central computation. However, when energy is a constraint and the environment is dynamic, distributed approaches are attractive and possibly are the only practical approach [20] because they avoid transmitting dynamic state information repeatedly across the network.

Previous work has presented adaptive distributed MAC and Transport level congestion response [13, 16]. Pottie and Kaiser [20] initiated work in the general area of wireless sensor networks by establishing that scalable wireless sensor networks require multi-hop operation to avoid sending large amounts of data over long distances. They went on to define techniques by which wireless nodes discover their neighbors and acquire synchronism. Given this basic bootstrapping capability, our work addresses the next level of automatic configuration that will be needed to realize envisioned sensor networks, namely, how to form the multi-hop topology [8]. Given the ability to send and receive packets, and the objective of forming an energy-efficient multi-hop network, we apply well-known techniques from MAC layer protocols to the problem of distributed topology formation. Similar techniques have also been applied to multicast transport protocol adjustment of periodic messaging [9].

In the following section we describe a sensor network scenario, stating our assumptions and considerations. Section 3 describes ASCENT in more detail. In Section 4, we present some initial experimental results using ASCENT. Related work is reviewed in Section 5.

Manuscript received Sep 02, 2014

K.Seena Naik, Department of CSE, S.K University, Anantapur, India
Dr. Ramachandra, Professor, Department of CSE, S.K University, Anantapur,India
M.V.Brahmananda Reddy, Department of CSE, GITAM University, Bangalore,India

II. Distributed Sensor Network Scenario

To motivate our research, consider a habitat monitoring sensor network that is to be deployed in a remote forest. Deployment of this network can be done, for example, by dropping a large number of sensor nodes from a plane, or placing them by hand. In this example, and in many other anticipated applications of ad-hoc wireless sensor networks, the deployed systems must be designed to operate under the following conditions and constraints:

- **Ad-hoc deployment**: we cannot expect the sensor field to be deployed in a regular fashion (e.g. a linear array, 2-dimensional lattice). More importantly, uniform deployment does not correspond to uniform connectivity owing to unpredictable propagation effects when nodes, and therefore antennae, are close to the ground and other surfaces, such us, obstacles, trees, etc.
- **Energy constraints**: The nodes (or at least some significant subset) will be untethered for power as well as communications and therefore the system must be designed to expend as little energy as is possible in order to maximize network lifetime.
- **Unattended operation under dynamics**: the anticipated number of elements in these systems will preclude manual configuration, and the environmental dynamics will preclude design-time pre-configuration.

In many such contexts it will be far easier to deploy larger numbers of nodes initially than to deploy additional nodes or additional energy reserves at a later date (similar to the economics of stringing cable for wired networks). In this paper we present one way in which nodes can exploit the resulting redundancy in order to extend system lifetime.

If we deploy too few nodes, the distance between neighboring nodes will be too great and the packet loss rate will increase; or the energy required to transmit the data over the longer distances will be prohibitive. If we use all deployed nodes simultaneously, the system will be expending unnecessary energy, at best, and at worst the nodes may interfere with one another by congesting the channel. In the process of finding an equilibrium, we are not trying to use a distributed localized algorithm to identify a single unique optimal solution. Rather this form of adaptive self-configuration using localized algorithms is well suited to problems spaces that have a large number of possible solutions. Our experimental results confirm that this is the case for our application.

In the next section, we describe a detailed algorithm for achieving our task. First we enumerate the following assumptions that apply to the remainder of our work since only a single implementation is studied and we do not have the flexibility afforded by a simulation.

- We assume a CSMA MAC protocol. This clearly introduces the possibilities for resource contention when too many neighboring nodes participate in the multihop network. Our basic approach should be relevant to TDMA MACs because distributed slot allocation schemes will also have degraded performance with increased load.
- Our algorithm reacts when links experience high packet loss. We do not detect or repair complete network partitions. We do not deal with network partitions because we address systems whose node density would preclude nodes from being completely isolated from one another. When node density is low, our approach is not needed because in general all nodes will be needed to form an effective network. Of course network partitions can occur in dense arrays when a swath of nodes are destroyed or obstructed. When such network partitions do occur, we believe that complementary system mechanisms will be most useful. For example, exploiting special resources such as long range radios deployed on a subset of deployed nodes, and used sparingly because of the power required. We leave such complementary techniques for network partition detection and repair to future work.
- ASCENT is not a routing or data dissemination protocol. ASCENT simply decides which nodes should join the routing infrastructure. Ad-hoc routing, Data Diffusion, or some other data dissemination mechanism, then runs over this multihop topology. In this respect, routing protocols are complementary to ASCENT.

The following section describes the ASCENT protocol in detail.

III. ASCENT

ASCENT consists of several phases. When a node first initializes, it enters into a listening-only phase called *neighbor discovery phase*, where each node obtains an estimate of the number of neighbors actively transmitting messages based on local measurements. Upon completion of this phase, nodes enter into the *join decision phase*, where they decide whether to join the multi-hop diffusion sensor network. During this phase, a node may temporarily join the network for a certain period of time to test whether it contributes to improved connectivity. If a node decides to join the network for a longer time, it enters into an *active phase* and starts sending routing control and data messages. If a node decides not to join the network, it enters into the *adaptive phase*, where it turns itself off for a period of time, or reduces its transmission range.

In this section, we describe these phases and their components. Several design choices present themselves in the context of ASCENT. We elaborate on these design choices while we describe the design it. Our initial experiments (Section 4) focus only on a subset of these design choices.

Introduction

Consider a simple sensor network for data gathering similar to the network described in Section 2. We cannot expect the sensor field to have uniform connectivity due to unpredictable propagation effects in the environment. Therefore, we would expect to find regions with low and high density. As we pointed out in Section 2, ASCENT does not deal with complete network partitions; we assume that there is a high enough node density to connect the entire region. Figure 1 shows a simplified schematic for ASCENT during initialization in a high-density region. For the sake of clarity, we are showing only the formation of a two-hop network. This analysis may be extended to networks of larger sizes.

Initially, only some nodes join the multi-hop network. The other nodes remain passively listening to the messages but not transmitting. This situation is depicted in Figure 1(a). The source starts transmitting data messages toward the sink. Because the sink is at the limit of radio range, it gets very high message loss from the source. The notion of what constitutes *high* is application dependent, and in our scheme, this is a parameter tunable by the application. We call this situation a *communication hole,* i.e. the receiver gets high message loss due to poor connectivity with the sender. The sink then starts sending help messages to signal neighbors that are in listen-only mode –also called *passive neighbors*– to join the network. It is called a help message because its goal is to ask for help to solve a communication problem due to lack of connectivity, i.e. help to repair a *communication hole*.

Join Decision Phase

The purpose of this phase is to decide whether to join the multi-hop sensor network. This is an important issue since the addition of a new node may change the total throughput of the network region (positively or negatively).

How can this determination be made and what variables should be used in the decision process? Several variables could be used to make a determination, such us, density of nodes, energy remaining, message loss, and events sensed. For each of these variables, we could use their absolute value, or even simpler, make them a binary function depending if their value is above or below a certain threshold.

Message Loss (ML)	**Number of Neighbors (NB)**	
	LOW	**HIGH**
LOW	If received help message, then temporarily joins	Do not join
HIGH	Join	Do not join

Table 1. Join Decision Engine.

As we mentioned previously, increase density of nodes produces more energy consumption in the contention for resources. Due to this, we have chosen density as one of our inputs in the decision process. One of our goals is to provide complete communication coverage to the sensor network. This cannot be achieved without a notion of *data message loss*[2]. This is the other local variable we have chosen as an input of our algorithm. In addition to these local variables, we include external signals triggered by neighbors. In our case,

help messages trigger a request to nodes in the vicinity, and we use the presence of such a message in our algorithm.

Upon entering in this phase, nodes already have an estimate of the number of active neighbors (NB) and the data message loss (ML). Table 3 summarizes the basic decision engine. In our particular implementation, each of these variables can be in two (2) states; it can be either HIGH or LOW, depending if they are above of below a certain threshold. When NB is HIGH, adding an additional node to the network will only contribute to cause more interference and contention for resources with the other nodes, even in the case when ML is LOW. Our goal is to maximize the lifetime of the network, providing only the minimum communication coverage necessary. This decision may need to be reviewed if we have a different goal in mind, such us providing multiple alternative paths, or improving sensing instead of communication coverage. The second case is when the NB is LOW and ML is HIGH. In this situation, we assume we are in the presence of a *communication hole,* and the node joins the network. An alternative we considered here was to join the network *only* in the presence of a help message. The problem with this alternative is that nodes at the edges of the network would never join, since there isn't any active downstream node sending help signals. The final case is when NB is LOW and ML is LOW. The node *only* temporarily joins if it receives a help message from some other node in the vicinity in order to probe the network. If the ML is still LOW after the probe phase period expired, then the node joins the multi-hop network and enters into the active phase. If the ML changes to HIGH after the probe phase period expired, then the node does not join the network and enters into the adaptive phase.

Active and Adaptive Phases

Nodes in the active phase send and receive all the control routing messages. When a node joins the network, it immediately sends a neighbor announcements message. This message signals other passive neighbors to increase their active neighbor count (and abort joining if necessary). When a node that has already joined the multi-hop network detects HIGH *data message loss* it sends help messages to other passive nodes.

If a node decides not to join the network, then the question is: what does it do? The answer to this question depends on the problem the sensor network is trying to solve. Nodes could turn themselves off for a certain period of time in order to save energy and extend the network lifetime. They could also reduce their radio range to reduce communication interference while maximizing sensing coverage. Since the goal of our original implementation was to maximize communication coverage and lifetime of the network, we picked the first form of adaptation for our experiments.

We emphasize that, even though we have discussed ASCENT's algorithms in some detail, much experimentation and evaluation of the various mechanisms and design choices is necessary before we fully understand the robustness, scale and performance of self-configuration. In particular, it is important to test our assumptions with real field experiments, since the level of dynamics we find in the real world is difficult to be captured completely in a simulator. The following section takes an initial step in this direction.

IV. Experimental Results

In this section, we report results from preliminary performance evaluation of ASCENT. We describe the experimental testbed used, our methodology, compare the performance of different variations of ASCENT's algorithms with a static version, and explore the impact of network dynamics and design parameters on our scheme.

Experimental Testbed

Before we proceed to the experimental results, we briefly describe the basic components of our testbed.

Figure 1 shows pictures of the hardware components of our testbed.

The basic node used in our testbed is the PC-104 node (Figure 2(a)). It consists of a basic computation platform based on the PC-104 stack [18], and an RPC radio [21] (Figure 2(c)). The DebugStations (Figure 2(b)) are similar to our PC-104 nodes, except that they have an Ethernet connector, and 8-port serial port card, and they do not have RPC radios. The PC-104 nodes connect to the DebugStations through ppp-serial connections, and from there to Ethernet. This connection is used as an out of band channel for logging and management purposes.

Figure 3 depicts the basic configuration of our testbed. After connecting the DebugStations to Ethernet and the PC-104 nodes to the DebugStation through a serial connection, the *entire* process runs automatically. PC- 104 nodes connect to the Master Debug Host at boot time to check if they need to update the software running. If they need it, they do the update and boot again, otherwise they start running the experiment with the user's software. All the experiments' traces are reliably logged on the Master Debug Host for post-processing. Multiple users can run experiments at the same time, as long as they do not interfere with each other (radio interference).

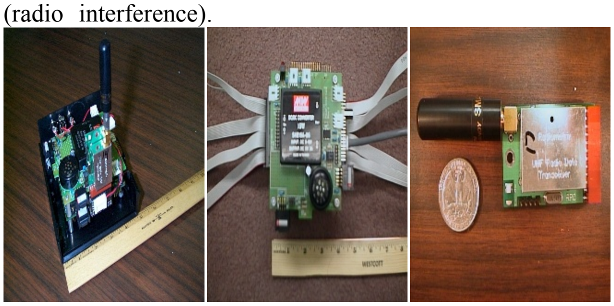

(a)PC-104 Node (b) PC-104 DebugStation (c) RPC Radiometrix Radio

Figure 1. Hardware components

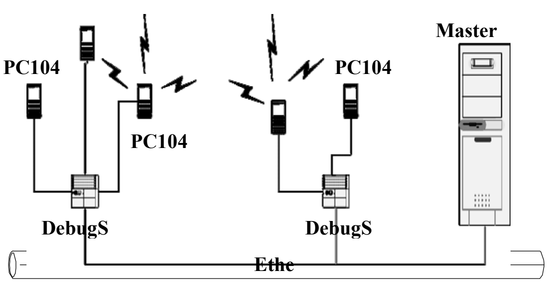

Figure 2. Basic Testbed configuration

Goals, Metrics, and Methodology

Our goals in evaluating ASCENT were four-fold: First, perform real experiments where we compare ASCENT with a statically configured scheme. This serves the important purpose of validating in reality some of our assumptions in designing the algorithms. Second, understand the impact of dynamics –such us node failures– on ASCENT. Third, explore the influence of randomization in ASCENT performance. Finally, study the sensitivity of ASCENT performance to the choice of parameters.

We choose 5 metrics to analyze the performance of ASCENT: Message Loss measures the percentage of messages not received by any node in the network. This metric indicates the real bandwidth available to the nodes in the sensor network. Control Overhead measures the ratio between control and data messages in the network. In sensor networks one of the most energy-consuming operations is communication. This metric represents the additional energy the network uses (by sending control extra messages) in order to be self-configurable. This metric includes ASCENT's help and neighbor announcement message, but does not include control messages sent by the underlying routing algorithm. Event Delivery Ratio is the ratio of the number of distinct event messages received by the sink to the number originally sent by the source. A similar metric has been used in ad-hoc routing [2]. Average Delay measures the average one-way latency observed between transmitting an event by the source and receiving it at the sink. This metric defines the latency added by our scheme. Average Reaction Time measures the average number of data message cycles since the occurrence of a failure until the sink receives the next data message. This metric measures the response time of the system when node failures occur.

In all our experiments we operate the sensor network far from overload. Hence, our sensor nodes do not experience congestion. Understanding the performance implications of congestion on our algorithms is the subject of future research. In spite of experimenting with uncongested networks, our nodes can incur in message losses, in particular due to dynamics and interference.

In order to study the performance of ASCENT's algorithms as a function of density, we run experiments ranging from 5 to 25 nodes in increments of 5 nodes. The experiments consist of 2 nodes placed near the maximum communication range from each other, and the rest of the nodes distributed between them. It is important to note that 25 nodes in this simple 2-hop network will lead to a very large network if the same density is used in every hop; a 2-d network of 10 hops by 10 hops would consist of roughly 2500 nodes. Moreover, we do not need to run larger multi-hop experiments to test our scheme since ASCENT does local self-configuration; all of our algorithms use only local neighbor information that comes from the next hop.

Each experiment sends approximately 100 messages per trial with one transmitter. Our results are averaged on the results of three trials. All the experiments were done in an indoors environment, with all the traditional obstacles you may find in an office-style kind of environment, such us, furniture, walls, doors, people, etc. The size of the messages used ranges from 100 to 150 bytes. The RPC maximum fragment size is 27 bytes, which implies that between 4 to 6 radio fragments are sent per message. The loss of a fragment implies the loss of the entire message, and we do not a have a MAC retransmission mechanism in our testbed. The RPC driver [22] implements a CSMA style MAC with simple collision avoidance mechanism. It will not emit packets until a (randomized) quiet interval has passed, i.e. an interval during which no packets have been received from other radios. We call this mandatory waiting time between receiving and sending the hold-off period. These simple scheme works because packet arrivals are usually correlated due to fragmentation. Multiple stations that are waiting to transmit will hopefully not collide due to the randomization of the hold-off period. In our experiments, the hold-off period was set to 100 ms and the randomization time to 200 ms. We set these values by carefully measuring the minimum time required by the RPC hardware and driver to transmit consecutive fragments. ASCENT also provides an application level randomization that is set to a maximum of 10 seconds. The data rate is set to 3 messages/min in most of the experiments. We use an implementation of data diffusion [12] as our routing protocol. Underlying routing messages were generated every minute, and the routing timeout was 3 min.

V. Related Work

Our work has been informed and influenced by a variety of other research efforts.

K. Sohrabi and G. Pottie [26] have made significant progress in self-configuration and synchronization in sensor networks at the single cluster level with a TDMA scheme. This work shares with us similar design principles, although it's more focused on low-level synchronization necessary for network self-assembly, while we concentrate on efficient multi-hop topology formation.

J. L. Gao's thesis [11] presented an adaptive local network formation/routing algorithm that facilitates cooperative signal processing. An election algorithm is used to select a central node among a small group of nodes that cooperate in information processing. While these algorithms were designed to operate for a relatively short time span in a reduced area near the target event, our objective is stable, long range topology formation that covers the entire sensor network.

The adaptive techniques we use were studied extensively to make the MAC layer self-configuring and adaptive more than 20 years ago during the refinement of contention protocols. More recently SRM [9] and RTCP [25] borrowed these techniques to adaptively adjust parameters such as session message frequency and randomization intervals. In this work we use those techniques to adapt the topology of a multi-hop wireless network.

Mobile ad-hoc networks [14, 17, 19] and Diffusion [12] adaptively configure the routing or data dissemination paths, but they do not adapt the basic topology. Qun Li and Daniela Rus [15] presented a scheme where mobile nodes modify their trajectory to transmit messages in the context of disconnected ad-hoc networks. This work shares with us the notion of adaptation of the basic topology for efficient delivery of messages, but it does so sending location updates between neighbors and using active messages to incrementally propagate them toward the destination. Our work uses measurements of neighbor density and message loss to exploit the redundancy of highly dense areas in the system in an energy efficient way. This work may complement ours in case of mobile nodes deployment and in the presence of network partitions. Ramanathan et. al. [23] proposed some distributed heuristics to adaptively adjust node transmit powers in response to topological changes caused by mobile nodes. This work assumes that a routing protocol is running at all times and provides basic neighbor information that is used to dynamically adjust transmit power. In our case, ASCENT decides which nodes should run the routing algorithm, and it makes this determination based on message loss in addition to density. Similarly, Ya Xu et. al. [27] also modify the topology in response to local measurements of density. Nevertheless, it uses geographic information, and it is not based strictly on measurements of message loss.

Self-configuration based on local measured parameters takes some inspiration from biological systems, in particular the models of ant colony behavior [4]. Bulusu et. al. [3], have proposed different algorithms for incremental beacon placement in sensor networks. This work share with us the same design principles, such us the use of localized algorithms, and adaptation based on locally measured parameters. While their work is oriented to solve the localization problem, ours is more oriented to energy efficient communication coverage.

Several centralized approaches have been developed to find optimal solutions for the localization, communication coverage, and other similar problems. Doherty et. al. [5] presented a method for estimating unknown node positions based exclusively on connectivity-induced constraints. These techniques could be adapted to find an optimal solution to the communication coverage problem based on location estimates. A novel theorem in combinatorial geometry that leads to a family of approximation algorithms for the geometric disk-covering problem has been proposed in [10]. It derives a centralized solution to find the minimum number of disk to cover certain points in the network. These points could be the nodes that experience high message loss, and the placement of the disks could represent the location of the nodes to be activated. This work provides very useful bounds to what is achievable but is not intended for application to energy-constrained contexts where the cost of extracting the information about node state to a central site is prohibitive. Moreover, the radio propagation characteristics may not be always best represented as disks of a certain radius. The work is being extended to distributed forms and may complement our work at a future date.

Conclusions and Future Work

In this paper, we described the design, implementation, and experimental evaluation of ASCENT, an adaptive self-configuration topology mechanism for distributed wireless sensor networks. There are many lessons we can draw from our preliminary experimentation. First, ASCENT has the potential for significant reduction of message loss and increase in energy efficiency. Even without optimizing parameter selection ASCENT outperforms the static scheme. Second, ASCENT mechanisms were responsive and stable under systematically varied conditions.

In the near future, we will perform experiments with larger numbers of nodes to further explore the scalability of our algorithms. We will also explore the use of wider area links to detect network partitions. In the longer term, we will explore more generally the limitations of localized algorithms and their relation to global information in the context of self-configurable sensor networks. We will also expand this work to address other modalities beyond communication coverage, such us, actuation or sensing.

This work is an initial foray into the design of self-configuring mechanisms for wireless sensor networks. Our distributed sensing network experiments represent a non-trivial exploration of the problem space. Such techniques will find increasing importance as the community seeks ways to exploit the redundancy offered by cheap, widely available microsensors, as a way of addressing new dimensions of network performance such as network-lifetime.

References

[1] B. Badrinath, M. Srivastava, K.Mills, J.Scholtz, and K.Sollins, Eds. Special Issue on Smart Spaces and Environments. IEEE Personal Communications, Oct. 2000.

[2] J. Broch, D. Maltz, D. Johnson, Y.-C. Hu, and J. Jetcheva. A Performance Comparison of Multi-Hop Wireless Ad-Hic Network Routing Protocols. In Proceedings of the Fourth Annual ACM/IEEE International Conference on Mobile computing and Networking (Mobicom'98), Dallas, TX, 1998.

[3] N. Bulusu, J. Heidemann and D. Estrin. AdaptiveBeaconPlacement. In Proceedings of the Twenty First International Conference on Distributed Computing Systems (ICDCS-21), Phoenix, Arizona, USA. April 16-19 2001. To Appear.

[4] G. Di Caro and M. Dorigo. AntNet: A Mobile Agents Approach to Adaptive Routing. Technical Report 97-12, IRIDIA, Universite' Libre de Bruxelles, 1997.

[5] L. Doherty, K. S. J. Pister, L. El Ghaoui. Convex Position Estimation in Wireless Sensor Networks. In Proceedings IEEE Infocom 2001, Anchorage AK, April 2001.

[6] J. Elson and D. Estrin. Random, Ephemeral Transaction Identifiers in Dynamic Sensor Networks. In Proceedings of the Twenty First International Conference on Distributed Computing Systems (ICDCS-21), Phoenix, Arizona, April 2001. To appear.

[7] D. Estrin, R. Govindan, and J. Heidemann, Eds. Special Issue on Embedding the Internet. Communications of the ACM, vol. 43, no. 5, May 2000. [8] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Next century challenges: Scalable Coordination in Sensor Networks. In Proceedings of the Fifth Annual International Conference on Mobile Computing and Networks (MobiCOM '99), August, 1999, Seattle, Washington

[8] S. Floyd, V. Jacobson, C-G. Liu, S. McCanne, and L. Zhang. A Reliable Multicast Framework for Lightweight Sessions and Application Level Framing. IEEE/ACM Transactions on Networking, November 1997.

[9] M. Franceschetti, M. Cook, and J. Bruck. A Geometric Theorem for Approximate Disk Covering Algorithms. Technical Report ETR035, Caltech University, January 2001.

[10] J. Gao. Energy Efficient Routing for Wireless Sensor Networks. PhD thesis in Electrical Engineering. University of California Los Angeles (UCLA). August 2000.

[11] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed Diffusion: A Scalable and Robust Communication Paradigm for Sensor Networks. In Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking (MobiCOM 2000), Boston, Massachusetts, August 2000

[12] V. Jacobson. Congestion Avoidance and Control. Proceedings of SIGCOMM '88, pages 314-329. Palo Alto, CA, August 1988.

[13] D. Johnson and D. Maltz. Dynamic Source Routing in Ad-hoc Wireless Networks. In T. Imielinski and H Korth, editors, Mobile Computing, pages 153-181. Kluwer Academic Publishers, 1996.

[14] Q. Li and D. Rus. Sending Messages to Mobile Users in Disconnected Ad-hoc Wireless Networks. In Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking (MobiCOM 2000), pages 44-55, Boston, August 2000.

[15] R. Metcalfe and D. Boggs. Ethernet: Distributed Packet Switching for Local Computer Networks. Communications of the ACM, 19 (5): 395-404, July 1976.

[16] [17] V. Park and M. Corson. A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks. In Proceedings of INFOCOM 97, pages 1405-1414, April 1997

[17] PC-104 Consortium, http://www.pc104.org

[18] C. Perkins and E. Royer. Ad hoc On-Demand Distance Vector Routing. Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications, New Orleans, LA, February 1999, pp. 90-100.

[19] G. Pottie and W. Kaiser. Wireless Integrated Network Sensors. Communications of the ACM, 43 (5): 51-58, May 2000. [21] Radio Packet Controller, http://www.radiometrix.com

[20] RPC Linux driver, http://www.circlemud.org/~jelson/software/radiometrix

[21] [23] S. Ramanathan and R. Rosales-Hain. Topology Control of Multihop Radio Networks using Transmit Power Adjustment. In Proceedings of IEEE Infocom 2000, Tel Aviv, Mar 2000.

[22] Sensors: The Journal of Applied Sensing Technology.

[23] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RFC 1889 RTP: A Transport Protocol for Real-Time Applications. Standards Track. January 1996.

[24] K. Sohrabi and G. Pottie. Performance of a Novel Self-Organization Protocol for Wireless Ad-Hoc Sensor Networks. In Proceedings of IEEE VTC, Amsterdam, Netherlands, September 1999.

[25] Y. Xu, J. Heidemann, and D. Estrin. Adaptive Energy-Conserving Routing for Multihop Ad Hoc Networks. Research Report 527, USC/Information Sciences Institute, October 2000.