

ASCENT: Adaptive Self-Configuring Sensor Networks Topologies

Alberto Cerpa and Deborah Estrin

{*cerpa,destrin*}@cs.ucla.edu

Center for Embedded Networked Sensing,
University of California, Los Angeles (UCLA)

Abstract

Advances in micro-sensor and radio technology will enable small but smart sensors to be deployed for a wide range of environmental monitoring applications. The low per-node cost will allow these wireless networks of sensors and actuators to be densely distributed. The nodes in these dense networks will coordinate to perform the distributed sensing and actuation tasks. Moreover, as described in this paper, the nodes can also coordinate to exploit the redundancy provided by high density, so as to extend overall system lifetime. The large number of nodes deployed in these systems will preclude manual configuration, and the environmental dynamics will preclude design-time pre-configuration. Therefore, nodes will have to self-configure to establish a topology that provides communication under stringent energy constraints.

ASCENT builds on the notion that as density increases, only a subset of the nodes are necessary to establish a routing forwarding backbone. In ASCENT, each node assesses its connectivity and adapts its participation in the multi-hop network topology based on the measured operating region. This paper motivates and describes the ASCENT algorithm and presents analysis, simulation and experimental measurements. We show that the system achieves linear increase in energy savings as a function of the density and the convergence time required in case of node failures while still providing adequate connectivity.

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. The sensor nodes will perform significant signal processing, computation, and network self-configuration to achieve scalable, robust and long-lived networks [2], [11], [10]. 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 simulation and experimental performance studies for a form of adaptive self-configuration designed for sensor networks. As we argue in Section II, 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 assesses 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 packet loss, requesting additional nodes in the region to join the network in order to relay messages.
- Reduces its duty cycle if it detects high packet 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 [27] because they avoid transmitting dynamic state information repeatedly across the network.

Pottie and Kaiser [27] 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 [10]. 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 been applied to multicast transport protocol adjustment of periodic messaging [13], [12].

In the following section we present a sensor network scenario, stating our assumptions and contributions. Related work is reviewed in section III. Section IV describes ASCENT in more detail. In Section V, we present some initial analysis, simulation and experimental results using ASCENT. Finally, in section VI we conclude.

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 [7], 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.
- 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 use too few of the deployed 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 optimal solution. Rather this form of adaptive self-configuration using localized algorithms is well suited to problem spaces that have a large number of possible solutions; in this context a large solution space translates into dense node deployment. Our simulation and experimental results confirm that this is the case for our application.

We enumerate the following assumptions that apply to the remainder of our work:

We assume a Carrier Sense Multiple Access (CSMA) MAC protocol with capacity to work in promiscuous mode. This clearly introduces the possibilities for resource contention when too many neighboring nodes participate in the multi-hop network. Our approach should be relevant to TDMA MACs as well because distributed slot allocation schemes will also have degraded performance with increased load. Future work will investigate the use of ASCENT with other MAC protocols under development [37].

Our algorithm reacts when links experience high packet loss. The ASCENT mechanism does not detect or repair network partitions of the underlying raw topology. Partitions are more prevalent when node density is low, and our approach is not applicable because in general all nodes will be needed to form an effective network. Of course network partitions can occur even in dense arrays when a swath of nodes are destroyed or obstructed. When such network partitions do occur, complementary system mechanisms will be needed; for example, detecting partitions in the multi-hop sensor network by exploiting information from long range radios deployed on a subset of nodes, and used sparingly because of the power required. We leave such complementary techniques for network partition detection and repair to future work.

The two primary contributions of our design are:

- The use of adaptive techniques that permit applications to *configure* the underlying topology based on their needs while trying to save energy to extend network lifetime. Our work does not presume a particular model of fairness, degree of connectivity, or capacity required.
- The use of self-configuring techniques that react to operating conditions *measured locally*. Our work is not restricted to the radio propagation model, the geographical distribution of nodes, or the routing mechanisms used.

III. RELATED WORK

Our work has been informed and influenced by a variety of other research efforts. There has been a great deal of work in the area of topology control, mostly using theoretical analysis or simulation, and involving MAC and power control mechanisms.

There have been several important theoretical evaluations of topology control. Most of this work focuses on the analysis of algorithms for distributed construction of a connected dominating set (CDS) of the corresponding unit-disk graph and the routing strategies using the CDS backbone [15], [34], [1], [16]. Gao et. al. [16] present a randomized algorithm for maintaining a CDS with low overhead. This algorithm assumes a randomized distribution of identifiers among all nodes, and the partition of the space in a grid, with at most one node selected as the cluster head in each grid. Gao's algorithm selects a small number of these cluster heads, and the total number selected has an approximation factor of $O(\sqrt{n})$ of the minimum theoretically possible. They also introduce a hierarchical algorithm for clustering and show an approximation factor of $O(1)$ with high probability. In later work, Gao et. al. present a distributed algorithm to construct a restricted Delaunay graph (RDG), where only Delaunay edges with a limited fix transmission radius are included [15]. This algorithm also uses the hierarchical clustering algorithm described previously [16]. The work shows that the number of edges in the restricted Delaunay graph is linear in the number of nodes, although the maximum degree of a node may be $\Omega(n)$ in the worst case. Alzoubi et al. [1] describe a distributed algorithm for constructing a minimum connected dominating set

(MCDS) for the unit-disk-graph with a constant approximation ratio of the minimum possible and linear time and message complexity. Wang and Li propose an algorithm to build a geometric spanner that can be implemented in a distributed manner [34]. The node degree is bounded by a positive constant, and the resulting backbone is a spanner for both hops and length.

The above algorithms provide an important theoretical foundation for topology control and help define the theoretical limits and bounds of what is achievable. Our work with ASCENT complements theirs by getting results from experiments using *real* radios, rather than using only simulation and analysis. Recent work [14], [6], [38], [35] evaluating radio connectivity using low-power radios suggests that these radio channels present *asymmetrical links*, *non-isotropic* connectivity, and *non-monotonic* distance decay of power with distance. It is important to understand the effects that these conditions impose on these topology control algorithms, since most of the real conditions observed using real radios violate the assumptions in the previous theoretical studies and may affect correctness. There is poor correlation between the spatial distance and reception rate, so assumptions based on geographic proximity between nodes do not necessarily hold in practice. Furthermore, the radio propagation is *not* circular, presenting non-isotropic properties. Finally, our previous work with SCALE [6] has shown the presence of asymmetric links for 5-30% of all pairwise communication, causing serious problems with algorithms that assume bidirectional connectivity.

The main approach followed by *MAC level protocols* to save energy has been to turn off the radios that do not have any scheduled transmission or reception of packets in a particular (usually small) timeframe. These protocols usually trade-off network delay for energy conservation because of the startup cost associated with turning the radios back on. K. Sohrabi and G. Pottie [32] 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. Sparse Topology and Energy Management (STEM) [31] accepts delays in path-setup time in exchange for energy savings. It uses a second radio (operating at a lower duty cycle) as a paging channel. When a node needs to send a packet, it pages the next node in the routing path. This node then turns on its main radio so that it can receive the packet. Sensor-MAC (S-MAC) [37] treats both per-node fairness and latency as secondary to energy conservation. It periodically turns off the radios of idle nodes and uses in-channel signaling to turn off radios that are not taking part in the current communication. More recent work [39] continue to explore MAC-level wake-up schemes. Most of the MAC schemes mentioned above are complementary to our work. ASCENT could establish a particular active topology and then use any of the above mechanisms to gain even further energy savings on the newly created active topology.

Another approach to reducing energy consumption has been to adaptively control the transmit power of the radio. The *lazy* scheduling proposed in Prabhakar et al. [28] transmits packets with the lowest possible transmit power for the longest possible time such that delay constraints are still met. Ramanathan et. al. [29] 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. While power control can be very useful, particularly in asymmetric networks such as cellular telephony, their advantages are less pronounced in sensor networks [6]. Furthermore, the power consumed by these low-power radios in idle state is of the same order of magnitude than the Tx or Rx state, so optimizations on transmit power are less important. Under these conditions, turning the radio off and putting the transceiver in sleep state is essential to extend network lifetime.

In Y. Xu et al. GAF [36], nodes use geographic location information to divide the network into fixed square grids. Nodes in each grid alternate between sleeping and listening, and there is always one node active to route packets per grid. ASCENT does not need any location aids, since it is based on connectivity. In addition, geographic proximity may not always lead to radio connectivity; this is why ASCENT uses local connectivity measurements. B. Chen et al. [8] proposed SPAN, an energy efficient algorithm for topology maintenance, where nodes decide whether to sleep or join the backbone based on connectivity

information supplied by a routing protocol. ASCENT does not depend on routing information, nor needs to modify the routing state; it decides whether to join the network or sleep based on measured local connectivity and packet loss information. In addition, our work does not presume a particular model of fairness or network capacity that the application requires.

J. L. Gao’s thesis [17] 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.

Mobile ad-hoc networks [21], [25], [26] and directed diffusion [19] adaptively configure the routing or data dissemination paths, but they do not adapt the basic topology. Q. Li and D. Rus [22] 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 by sending location updates between neighbors and using active messages to incrementally propagate them toward the destination. Our work uses measurements of neighbor density and packet loss to exploit the redundancy of 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.

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 [20], [23]. More recently SRM [13] and RTCP [30] 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.

Self-configuration based on local measured parameters takes some inspiration from biological systems, in particular the models of ant colony behavior [5]. Bulusu et. al. [4], have proposed different algorithms for incremental beacon placement in sensor networks. This work share with us the same design principles, such as the use of localized algorithms, and adaptation based on locally measured parameters. While their work is oriented to solve the localization problem, ours is oriented to energy efficient communication and sensing coverage.

The following section describes the ASCENT protocol in some detail.

IV. ASCENT DESIGN

ASCENT adaptively elects “active” nodes from all nodes in the network. Active nodes stay awake all the time and perform multi-hop packet routing, while the rest of the nodes remain “passive” and periodically check if they should become active.

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 of the underlying raw topology; we assume that there is a high enough node density to connect the entire region. Figure IV shows a simplified schematic for ASCENT during initialization in a high-density region. For the sake of clarity, we show only the formation of a two-hop network. This analysis may be extended to networks of larger sizes.

Initially, only some nodes are active. The other nodes remain passively listening to packets but not transmitting. This situation is depicted in Figure 1(a). The source starts transmitting data packets toward the sink. Because the sink is at the limit of radio range, it gets very high packet loss from the source. We call this situation a *communication hole*; the receiver gets high packet 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.

When a neighbor receives a *help message*, it may decide to join the network. This situation is illustrated in 1(b). When a node joins the network it starts transmitting and receiving packets, i.e. it becomes an *active*




Fig. 1. Network self-configuration

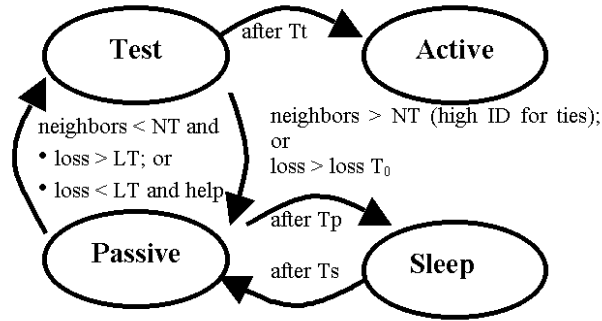


Fig. 2. ASCENT state transitions

neighbor. As soon as a node decides to join the network, it signals the existence of a new active neighbor to other passive neighbors by sending an *neighbor announcement message*. This situation continues until the number of active nodes stabilizes on a certain value and the cycle stops (see Figure 1(c)). When the process completes, the group of newly active neighbors that have joined the network make the delivery of data from source to sink more reliable. The process will re-start when some future network event (e.g. node failure) or environmental effect (e.g. new obstacle) causes packet loss again.

In this section, we describe the ASCENT algorithm and their components. Several design choices present themselves in this context. We elaborate on these design choices while we describe the design. Our initial analysis, simulations, and experiments in Section V focus only on a subset of these design choices.

A. ASCENT state transitions

In ASCENT, nodes are in one of four states: *sleep*, *passive*, *test*, and *active*. Figure 2 shows a state transition diagram.

Initially, a random timer turns on the nodes to avoid synchronization. When a node starts, it initializes in the *test state*. Nodes in the *test state* exchange data and routing control messages. In addition, when a node enters the *test state*, it sets up a timer T_t , and sends *neighbor announcement messages*. When T_t expires, the node enters the *active state*. If before T_t expires the number of active neighbors is above the *neighbor threshold* (NT), or if the average *data loss rate* (DL) is higher than the average loss before entering in the *test state*, then the node moves into the *passive state*. If multiple nodes make a transition

to the *test state*, then we use the node ID in the announcement message as a tie breaking mechanism (higher IDs win). The intuition behind the *test state* is to probe the network to see if the addition of a new node may actually improve connectivity.

When a node enters the *passive state*, it sets up a timer T_p and sends *new passive node announcement messages*. This information is used by active nodes to make an estimate of the total density of nodes in the neighborhood. Active nodes transmit this density estimate to any new passive node in the neighborhood. When T_p expires, the node enters the sleep state. If before T_p expires the number of neighbors is below NT , and either the DL is higher than the *loss threshold* (LT) or DL is below the *loss threshold* but the node received a *help message* from an active neighbor, it makes a transition to the *test state*. While in *passive state* nodes have their radio on, and are able to overhear all packets transmitted by their active neighbors (even if the packets are not addressed to the passive node, since the radio is in promiscuous mode). No routing or data packets are forwarded in this state, since this is a listen-only state. The intuition behind the *passive state* is to gather information regarding the state of the network without causing interference with the other nodes. Nodes in the *passive* and *test states* continuously update the number of active neighbors and data loss rate values. Energy is still consumed in the *passive state*, since the radio is still on when not receiving packets. A node that enters the *sleep state* turns the radio off, sets a timer T_s and goes to sleep. When T_s expires, the node moves into *passive state*. Finally, a node in *active state* continues forwarding data and routing packets until it runs out of energy. If the *data loss rate* is greater than LT , the active node sends *help messages*.

B. ASCENT parameters tuning

ASCENT has some parameters that can affect its final behavior. In this section, we explain the choices made in the current ASCENT algorithm. A particular application may select different values for some parameters, for instance, trading energy savings for greater reaction time in case of dynamics. ASCENT also provides adaptive mechanisms for the optimal determination of some parameter values dynamically at running time.

The *neighbor threshold* (NT) value determines the average degree of connectivity of the network. An application could adjust this value dynamically depending on the events occurring in a certain area of the network, for example, to increase network capacity. In this study, we set this value to 4.

The *loss threshold* (LT) determines the maximum amount of data loss an application can tolerate before it requests help to improve network connectivity. This value is very application dependent. For example, average temperature measurements from a sector of a forest will not tend to vary drastically, and the application may tolerate high packet loss. In contrast, tracking of a moving target by the sensor network may be more sensitive to packet losses. In our implementation this value was set to 20%.

The test timer T_t and the passive timer T_p determine the maximum time a node remains in the test and passive states, respectively. They face a similar trade-off of power consumption vs. decision quality. The larger the timers, the more robust the decision in presence of transient packet losses (that also affects the neighbor determination), but the greater the power consumed with the radio on, and viceversa. Our work with SCALE [6] has shown that the final determination of these timer values should be dependent on the quality of the reception rate for each link. On the one hand, links that present very high ($> 80\%$) or very low ($< 20\%$) reception rates show less variability over time, and consequently require less time to make an accurate determination of the link quality. On the other hand, links with intermediate reception rates show great variability over time and require more time to make better estimations. We note that it should be possible to design a mechanism that automatically determines the minimum amount of time we should measure the channel to provide some statistical bounds on the accuracy of the link quality estimation, but we have left this as future work. In our implementation, the T_p timer was set to 2 minutes and T_t to 4 minutes.

Similarly, the sleep timer T_s represents the amount of time the node sleeps to preserve energy. The larger the T_s timer, the larger the energy savings, but the larger also the probability of no node in passive

state ready to react to dynamics. ASCENT uses an adaptive probabilistic mechanism in order to determine the optimal relationship between T_p and T_s timers. This mechanism is solely dependent on the average density neighborhood estimate in the and the probability threshold P_t that a certain k number of nodes in the neighborhood are in passive state at any given point in time. The details of this mechanism are explained in section V-B. In our implementation, the value of k was set to 2 and P_t was set to 95%.

C. Neighbor and Data Loss determination

The number of active neighbors and the average data loss rate are values measured *locally* by each node while in passive and test state.

We have chosen to define a neighbor as a node from which we receive a certain percentage of packets over time. This implies having a history window function (CW) that keeps track of the packets received from each individual node over a certain period (time and/or number of messages), and a fixed or dynamic *neighbor loss threshold* (NLS).

In ASCENT, each node adds a unitary monotonically increasing sequence number to each packet transmitted (including data and control packets). This permits neighbor link loss detection when a sequence number is skipped. In addition, we assume application data packets also have some mechanism to detect losses (data payload sequence numbers in our implementation). Additionally, the final packet loss (or its reciprocal reception rate) estimate from each neighbor node is calculated by using an exponentially weighted moving average (EWMA) of the form:

$$EWMA_{current} = \rho \cdot CW + (1 - \rho)EWMA_{previous}$$

The value of the filter constant ρ was set to $0.\hat{3}$, which effectively means that only the last 5 estimates have any significant weight on the current estimate. This estimate provides a measure of the incoming packet loss from any node toward the neighbor doing the calculation, but it does not provide information of the packet loss perceived by the neighbors. In order to close the loop, this information is periodically exchanged between active and test nodes (not passive nodes) by piggybacking this information in data packets or by sending hello packets in the absence of data traffic.

The number of active neighbors N is defined as the number of neighbors with link packet loss smaller than the neighbor loss threshold (NLS) and with symmetrical links. In our study, we consider a link symmetrical if has a difference in reception rate of less than 40% between the incoming and outgoing reception rate. We have chosen the following formula NLS :

$$NLS = 1 - \frac{1}{N}$$

with N being the number of neighbors calculated in the previous cycle.

When a node gets a neighbor's packet loss estimate larger than the NLS , it no longer considers that node as a neighbor and deletes it from its neighbor list. The intuition behind this formula is the following: as we increase the number of neighbors in the region, the likelihood of any pair of them not listening to each other (or having high losses) increases. Therefore, as we increase the number of neighbors, we should correspondingly increase the neighbor's loss threshold. Not doing so may result in getting a lower neighbor count even though nodes in the region may still interfere with each other. Correspondingly, as we decrease the number of neighbors, we should decrease the neighbor's loss threshold accordingly. (We experimented with some other functions, like an inversely decaying function of $1/N$ and an exponentially decaying function of $1/N$; but the simple formula above worked best).

The average *data loss rate* (DL) is calculated based on the application data packets. Data losses are detected using data sequence numbers. Depending on the routing strategy, a node may receive multiple copies of the same application data packet. We only consider a data loss if the message was not received from any neighbor during a certain configurable period of time (this allows out of order delivery based on the application needs). Control messages (help, neighbor announcements and routing) are not considered in this calculation.

D. ASCENT interactions with routing

ASCENT runs above the link and MAC layer and below the routing layer. ASCENT is not a routing or data dissemination protocol. ASCENT simply decides which nodes should join the routing infrastructure. Ad-hoc routing [15, 18, 20], Directed Diffusion [13], or some other data dissemination mechanism, then runs over this multi-hop topology. In this respect, routing protocols are complementary to ASCENT.

ASCENT nodes become active or passive independent of the routing protocol running on the node. In addition, ASCENT does not use state gathered by the routing protocol, since this state may vary greatly for different protocols (e.g. ad-hoc routing tables and directed diffusion gradients), or requires changing the routing state in any way. Currently, if a node is testing the network and it is actively routing packets when it becomes passive, ASCENT depends on the routing protocol to quickly re-route traffic. This may cause some packet loss, and therefore an improvement that has not been implemented is to inform the routing protocol of ASCENT's state changes so traffic could be re-routed in advance.

We emphasize that, even though we have discussed the ASCENT algorithm 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. The following section presents our initial findings based on simple analysis, simulation, and an experimental implementation.

V. PERFORMANCE EVALUATION

In this section, we report results from a preliminary performance evaluation of ASCENT. We use simple mathematical models to determine an idealized expected performance of delivery rate, latency, and energy savings as we increase node density. Since our analysis cannot capture the complexity of a full ASCENT scenario, we use simulations and real experiments to further validate the performance evaluation.

A. Goals and Metrics

Our goals in evaluating ASCENT were three-fold: First, in order to validate some of the assumptions made during design of the algorithm, perform analysis, simulations and real experiments and conduct comparative performance evaluation of the system with and without ASCENT. Second, understand the energy savings and delivery rate improvements that can be obtained by using ASCENT. Finally, study the sensitivity of ASCENT performance to the choice of parameters.

We choose four metrics to analyze the performance of ASCENT: *One-Hop Delivery Rate* measures the percentage of packets received by any node in the network. When all the nodes are turned on –we call this the *Active case*– the packet reception includes all nodes. In the ASCENT case, it includes all nodes but the ones in the sleep state. This metric indicates the effective one-hop bandwidth available to the nodes in the sensor network. *End-to-End Delivery Rate* is the ratio of the number of distinct packets received by the destination to the number originally sent by the source. It provides an idea of the quality of the paths in the network, and the effective multi-hop bandwidth. A similar metric has been used in ad-hoc routing [3]. *Energy Savings* is the ratio of the energy consumed by the Active case to the energy consumed by the ASCENT case. This metric defines the amount of energy savings and network lifetime we gain by using the ASCENT algorithm. Finally, *Average Per-Hop Latency* measures the average delay in packet forwarding in a multi-hop network. It provides an estimate of the end-to-end delay in packet forwarding.

B. Analytic performance analysis

To understand the relationship between expected packet delivery and density of nodes we first use a simple mathematical analysis.

Assume that nodes are randomly distributed in an area A , and they have an average degree of connectivity of n . Further assume packets are propagated using flooding with a random back-off upon packet reception. This random component is chosen from a discrete pool of S slots with a uniform probability

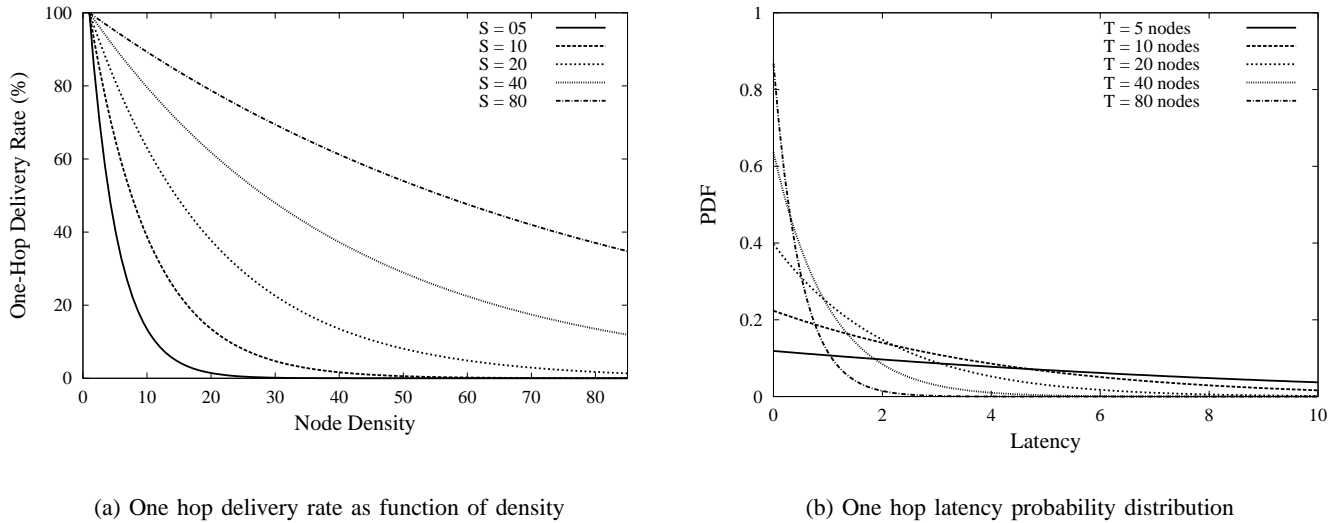


Fig. 3. Figure 3(a) shows the expected one hop delivery rate as a function of density. The larger the randomization period, the better the one hop delivery rate for any given density. Figure 3(b) shows the probability distribution of the one hop latency. The larger the density, the smaller the probability of a large latency.

distribution. Thus, the probability of successfully transmitting a packet with no collisions when there are T potential forwarding nodes in the vicinity is given by:

$$P(\text{success}) = \left(\frac{S-1}{S} \right)^T \quad (1)$$

From this formula we see that as we increase the density of transmitting nodes T , the probability of successfully delivering packets without collisions decreases proportionally. When all the nodes in the network are able to transmit and receive packets, we find that $T = n$, since every node in the vicinity can transmit packets (assuming a lossless channel, all nodes received the original packet). Increasing the density of nodes increases the probability of collisions in the area. ASCENT fixes the number of transmitters in the area to the *neighbor threshold* (NT) value, resulting in $T = NT$, independent of the total number of nodes, n , deployed. Figure 3(a) shows the analytical relation between expected one hop delivery rate vs. density of nodes for different S values.

The relation between the hop-by-hop latency introduced by the randomization and the density of nodes can be analyzed similarly. The average latency experienced per hop is related to the number of random slots S and the total number of active nodes T . After reception of a message to be forwarded toward the destination, each of the T active nodes picks a random slot, say S_1, S_2, \dots, S_T . The mean number of all the random slots chosen will tend to be $S/2$, since it is a uniform probability distribution. Assuming no collision losses, i.e.: $\forall i \neq j \in 1, 2, \dots, T \Rightarrow S_i \neq S_j$, the hop-by-hop latency is determined by the first message to be forwarded. The delay δ is then:

$$\delta = \min(S_1, S_2, \dots, S_T)$$

We want to find $P(\delta)$, the probability distribution of the smaller random time slot picked by T nodes. We define:

$$Q(y) = \text{Prob}[\min(S_1, S_2, \dots, S_T) > y] \Leftrightarrow \text{each of } (S_1, S_2, \dots, S_T) > y$$

This happens with probability:

$$\left(\frac{S-y}{S}\right)^T \therefore Q(y) = \left(1 - \frac{y}{S}\right)^T$$

$P(\delta)$ as defined above is:

$$P(\delta) = Q(\delta) - Q(\delta + 1) = \left(1 - \frac{\delta}{S}\right)^T - \left(1 - \frac{(\delta + 1)}{S}\right)^T \quad (2)$$

Figure 3(b) shows the $P(\delta)$ distribution for different values of T and $S = 20$. When all the nodes in the network are able to transmit and receive packets, we find that $T = n$. As n increases, the mean value of $P(\delta)$ decreases. This result corresponds to the intuition that as we increase the total number of transmitting nodes, the likelihood of any of them picking a smaller random value increases. In the ASCENT case, $T = NT$ independently of the density n , and the mean value of $P(\delta)$ remains constant.

Finally, we would like to understand the energy savings that could be obtained by using ASCENT. When the system is not running ASCENT, all the nodes have their radios on, consuming *Idle* power¹. When the system is running ASCENT, NT nodes have their radios on, while the rest alternate between sleeping and listening. The energy savings (ES) are:

$$ES = \frac{n \cdot Idle}{NT \cdot Idle + (n - NT) \cdot Idle \cdot \frac{T_p}{T_p + T_s} + (n - NT) \cdot Sleep \cdot \frac{T_s}{T_p + T_s}} \quad (3)$$

The numerator represents the power consumed by all the nodes when not running ASCENT. The denominator represents the power consumed by all nodes running ASCENT. The first term in the denominator indicates the power consumed by the NT nodes selected by ASCENT to have their radios on. The second term in the denominator indicates the energy of non-active nodes when in passive state, and the third term indicates the energy consumed while in sleep state. We define α to be the ratio of the passive timer T_p to the sleep timer T_s . We also define β to be the ratio of the radio's sleep mode to the idle mode power consumption. By replacing these new definitions in equation 3 we get:

$$ES = \frac{n}{NT + (n - NT) \cdot \frac{\alpha + \beta}{\alpha + 1}} \quad (4)$$

Equation 5 shows the upper bound of the energy savings as we increase density.

$$\lim_{n \rightarrow \infty} ES = \frac{\alpha + 1}{\alpha + \beta} \quad (5)$$

Figure 4 shows the energy savings as we increase the density of nodes for a fixed value of β . For a fixed NT value and a small value of β , as we increase density the power consumption is dominated by the passive nodes in the passive-sleep cycle. The intuition is that the smaller the α , the larger T_s in relation to T_p , and consequently, the larger the energy savings the system can achieve. Note that these savings come at a cost; the larger the T_s , the larger the reaction time of the system in case of dynamics. There is a trade off between the number of nodes we would like in passive state ready to react to dynamics and the energy savings we can achieve by having a more aggressive sleeping schedule. Even if we set up the network with an "optimal" value of α at initialization, the density of the network will not be homogeneous, and even if it is homogeneous initially, it will probably change as nodes drain their batteries and die out. To cope with this trade off, we propose an adaptive probabilistic mechanism where the value of α depends on the minimum probability of k nodes being in passive state at any given moment in time, and the density of nodes in the neighboring region.

If we assume that nodes alternate between passive and sleep state and that their schedules are indepen-

¹The difference in power consumption between the Tx, Rx, and Idle radio state is not significant. See Section V-C

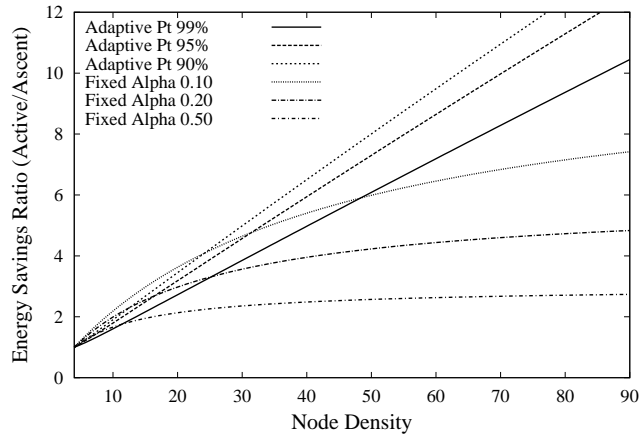


Fig. 4. Energy savings as a function of density for ASCENT fixed and adaptive state timers. Fixed state timers converge asymptotically to a particular maximum value when increasing density. On the other hand, adaptive state timers do not present this asymptotic limit and the energy savings increase linearly as a function of density.

dent, the probability of any node being in passive or sleep state is given by:

$$P(\text{passive}) = \frac{\alpha}{\alpha + 1} \quad P(\text{sleep}) = \frac{1}{\alpha + 1} \quad (6)$$

We want to find the minimum value of α such that the probability of at least k nodes being passive at any given moment in time is larger than a minimum probability threshold P_t . We call this value $\alpha_{P_t k}$. Any α value smaller than $\alpha_{P_t k}$ will not comply with the minimum probability requirement, and any α value larger than $\alpha_{P_t k}$ will comply with the minimum probability requirement but will expend unnecessary energy.

For the given state probabilities given in equation 6, the probability of at least k nodes in passive state at the same time is given by:

$$P(k) = 1 - \left(\frac{1}{\alpha + 1} \right)^n \cdot \frac{\alpha^k - 1}{\alpha - 1} \quad (7)$$

The Appendix shows the proof of equation 7. We want to find the values of α for different values of k being $P(k)$ equal to the minimum probability threshold P_t . For $k = 1$ (at least 1 passive node at any given time) and $k = 2$ (at least 2 passive nodes at any given time) there are closed-form solutions to the value of α ($\alpha_{P_t 1}$ and $\alpha_{P_t 2}$):

$$\alpha_{P_t 1} = 10^{-\frac{1}{n} \cdot \log(1-P_t)} - 1 \quad (8)$$

$$\alpha_{P_t 2} = 10^{\frac{1}{1-n} \cdot \log(1-P_t)} - 1 \quad (9)$$

For other values of k , the optimal value $\alpha_{P_t k}$ for a specific P_t and density n can be found by using iterative numerical techniques for the solution of nonlinear equations (e.g. Newton's method). The initial searching value x_0 could be set to $\alpha_{P_t 2}$. Figure 5(a) shows the optimal values of α for $k = 1$ and $k = 2$ for different probability thresholds.

Figure 4 shows the energy savings of ASCENT with adaptive state timers. In this case, we no longer have an asymptotic behavior as density increases like in the previous fixed timers case. The energy savings increase linearly with density, and the slope of the line is primarily determined by the probability threshold P_t (see below). Figure 5(b) shows that the impact of additional redundancy by incrementing the value of k (concurrent passive nodes) has less of an impact, and only reduces the energy savings by a minimal factor.

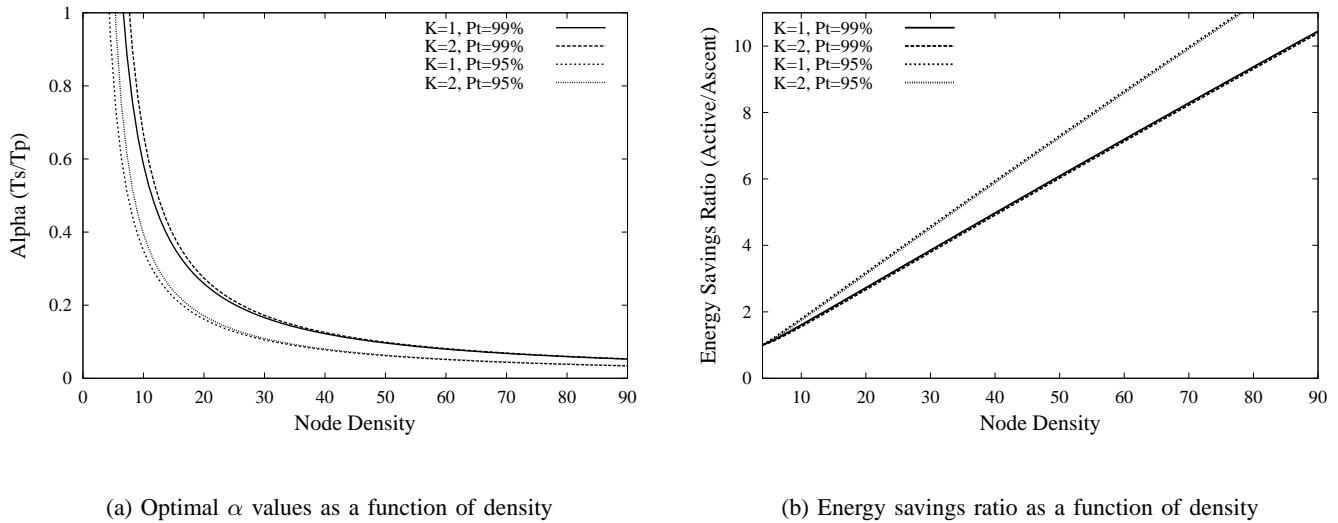


Fig. 5. Figure 5(a) shows the optimal α values for $k = 1$ and $k = 2$ for different probability thresholds. Figure 5(b) shows the energy savings ratio as a function of density for different values of k and P_t . It is clear from the graphs that the probability threshold P_t has the most noticeable effect in the determination of the value of α and the energy savings ratio.

C. Simulation and Experimental Methodology

Implementation: ASCENT implementation was developed using the EmStar programming environment [9]. We implemented ASCENT using a number of fine-grained modules, so that other developers could re-use as much of our work as possible. Figure 6 shows the diagram of the code structure. The first is the *LinkStats* module, which adds a monotonically increasing sequence number to each packet sent by any process on the node. It monitors such packets arriving from other nodes, and maintains detailed packet statistics for high precision connectivity measurements without increasing channel use (but, slightly reducing the maximum data payload). This module also implements the exponentially weighted moving average (EWMA) filter for the reception rate of each neighbor. The second module is *Neighbor Discovery*, which sends and receives heartbeat messages, and maintains a list of active neighbors. Third, to evaluate energy usage, we created the *Energy Manager* module that acts as a simulated battery for each node. It counts packets sent and received, idle time, and radios powering on and off; energy is deducted from an initial supply accordingly. It runs in two modes: either it merely tracks energy usage or it actually shuts off a node when that node runs out of energy. Finally, we created the *ASCENT* protocol implementation itself, which uses the information provided by the other modules.

Simulator: ASCENT was simulated using the built in simulator (*emsim*) provided by EmStar [9]. The simulator essentially runs exactly the *same* code base than the implementation, with no modifications. As in reality, the nodes must interact using their radios and are not allowed to share state directly. Instead of using real radios and sensors, *emsim* provides a channel simulator that models the behavior of the environment. The channel model uses information provided at configuration time, such as node position and transmit power. The channel model used in our simulations is a statistical model based on extensive radio connectivity traces gathered when developing earlier versions of ASCENT and SCALE [6]. The simulator is also able to provide CSMA style of collisions. *emsim* does not propagate packets further in the stack if two or more packets are received within the contention period. Thus, the probability of collisions is determined by the length of the contention period and the traffic.

Experimental Testbed: Figure 7 shows pictures of the hardware components of our testbed. The ceiling array [9] used in the experiments is composed of various serial port multiplexors attached to a testbed PC. Figure 7(b) shows an image of the ceiling array deployed in our lab. We use UTP Cat 5 cables of

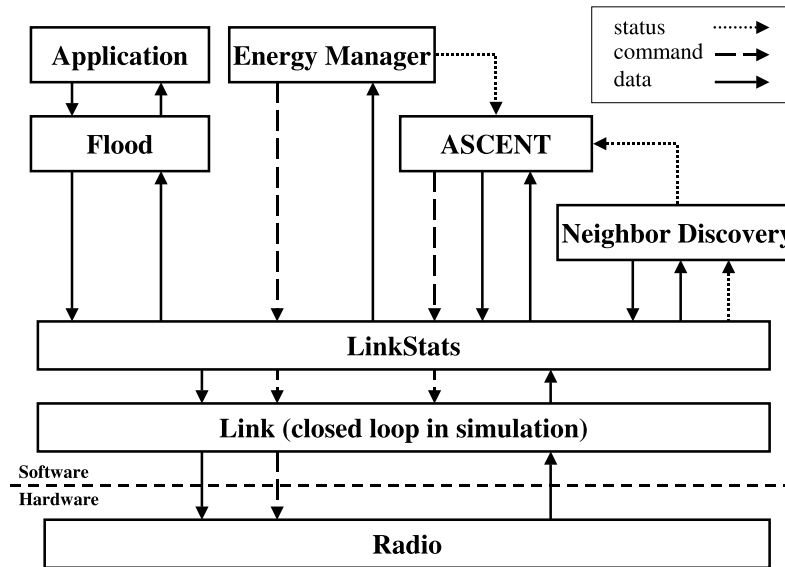


Fig. 6. ASCENT code structure. ASCENT was developed in a modular way, so other developers could re-use as much functionality as possible even when not running the ASCENT topology control algorithm.



(a) Mica 1 mote



(b) Indoor Office, UCLA CENS lab ceiling array

Fig. 7. Experimental Testbed. The ceiling array is composed of a PC attached to various serial multiplexors. Several UTP cables run from each multiplexor to the deployment locations in the ceiling where a mote is attached at the end.

different lengths (up to 30 meters) and attach on end of the cable to the multiplexor and the other end to a node. A total of 55 nodes are used in the testbed. The nodes are wall powered.

Figure 7(a) shows a picture of the Mica 1, the node used in the experimental testbed. Table I shows the main features of the hardware platform used. The Mica 1 mote firmware comes with an event-driven operating system called TinyOS [18]. It provides a DC-balanced single-error correction and double bit error detection (SECDED) scheme to encode each byte transmitted by the RF transceiver (RFM). The system supports variable packet sizes, and uses a 16-bit CRC that is computed over the entire packet for error detections. A simple driver (*Transceiver*) was used to run on the motes in TinyOS. Its function is to send/receive packets to/from the radio and pass them from/to the PC using a host-mote protocol over the serial connection.

TABLE I
NODE CHARACTERISTICS

	Mica 1
CPU Processor	Amtel 128
Prog. Memory (KB)	128
Data Memory (KB)	4
Serial RS232	needs adapter
Clock Speed (MHZ)	4
RF Manufacturer	RFM [24]
RF Transceiver	TR1000
Radio frequency (MHz)	916
Modulation	ASK
Throughput (kbps)	13.3
TX power [0dBm] (mW)	< 1
Antenna	Omni whip

Scenarios and environment: In order to study the performance of ASCENT’s algorithms as a function of density, we run experiments with different densities ranging from 5 to 40 nodes. In this study, density is defined topologically, i.e. the density of nodes is defined by the average degree of connectivity of all the nodes in the experiment and *not* by their physical location (geographical density). Since we could not easily change the location of nodes in the ceiling array, and since the physical size of our lab is limited, we achieved different levels of density by adjusting the transmit power of the RF transceiver. Using SCALE [6], we built the entire connectivity map of the ceiling array for different transmission power levels and picked the power level that provided us with the desired density. The average number of hops in the topologies obtained by this method was three. All the experiments were done in an indoors environment, with obstacles such as, furniture, walls, cubicles, doors, etc. The simulations replicate the same scenarios tried in the experiments. For each simulation, we vary the density of nodes from 5 to 80 nodes. In addition, for larger multi-hop simulations, we incremented the number of sources and destinations from 1 to 5. The average number of hops in the simulations was six. In all the experiments and simulations, the source(s) and the destination(s) were placed at the edge of the network to maximize the number of hops and usage of transit nodes (nodes transmitting traffic from/to the source/destination). Each experimental point in the graphs presented in the following sections is the average of three experimental trials, and each simulation point in the graphs represents the average of five simulated trials. All the results include confidence intervals with a degree of confidence of 95%.

Traffic: In each experiment, one source sends approximately 200 messages with temperature and light sensor readings (the readings were stored values). The data rate was set to 3 sensor reading messages per minute. In each simulation, one or more sources send approximately 400 messages each. The data rate was the same as the experiments. In all our experiments and simulations 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 packet losses due to dynamics and interference.

Routing: We use flooding as our routing protocol. In order to reduce contention when multiple nodes try to re-forward packets received at the same time, the flooding module has a programmable randomization interval. Upon receiving a packet, the flood module will wait for a random time between zero and the maximum randomization interval. In our experiments the randomization interval was set to 5 seconds (unless otherwise noticed).

Energy Model: To model the energy consumption, we looked at the manual specifications of the

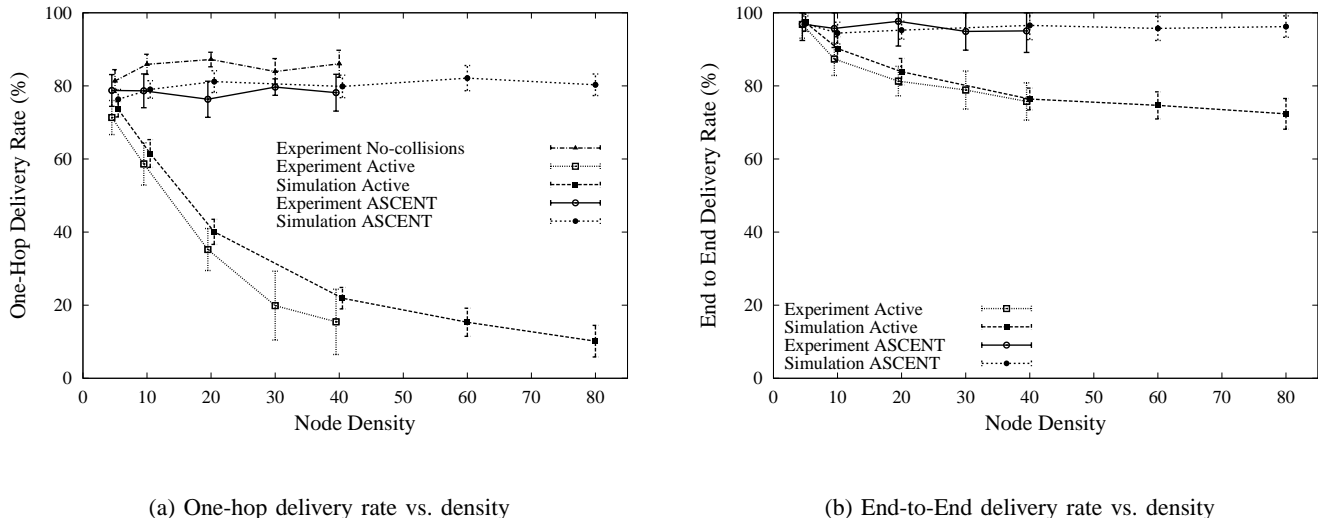


Fig. 8. Figure 8(a) shows the one-hop delivery rate as a function of density. ASCENT limits the number of active nodes forwarding traffic to NT and reduces contention for the channel. Figure 8(b) shows the end-to-end delivery rate as a function of density. ASCENT end-to-end delivery rate is stable for the range of densities tested.

RFM Tx-1000 [24]. We found that the values for Tx:Rx:Idle:Sleep in mW 36:9:9:0.015 for the RFM. Several studies [8], [33] have reported differences of the order of 10:1 between Idle and Sleeping power consumption for 802.11 wireless LAN cards. For the low power radios we study, this difference is in the order of 100:1. This relation is important since it is the β factor defined in the previous section. In our model, we did not consider the energy consumed by the CPU.

The remainder of this section presents our simulation and experimental results.

D. Network Capacity

Our first simulations and experiments compare the one-hop delivery rate and the end-to-end delivery rate of the system with and without ASCENT (with adaptive timers enabled).

Figure 8(a) shows the one-hop delivery rate as a function of the density in a multi-hop network. The “No-collisions” curve shows the average one-hop delivery rate in the network for the different densities tested. It shows the average losses due to environmental effects in the absence of simultaneous transmissions. The values were obtained by running SCALE [6] on the ceiling array testbed with different transmission power values. The results are encouraging. To a first degree, there are no important differences between the expected analytical and simulated performance and the performance using real radios up to densities of 40 nodes. In the Active case (no self-configuration, all nodes are turned on), all the nodes join the network and forward packets. This case has low delivery rate because as we increase the density of nodes, the probability of collisions increases accordingly when using flooding as a routing strategy. It rapidly reaches around 40% with densities of 20 nodes, and enters into a saturation region after that. ASCENT limits the number of active nodes to the NT value, and therefore does not increase channel contention with larger densities.

Figure 8(b) shows the end-to-end delivery rate, i.e. the percentage of packets transmitted by the source that reached the destination. In the experiments, each packet traverses an average of 3 hops. The simulations were done on a larger network, with packets traversing on average 6 hops from source to destination. We can see that ASCENT outperforms the Active case. ASCENT’s performance remains stable as the density increases, which demonstrates the scalability properties of our algorithms as the density increases. The Active case does not perform as bad as one would expect based on the one-hop delivery rate shown in

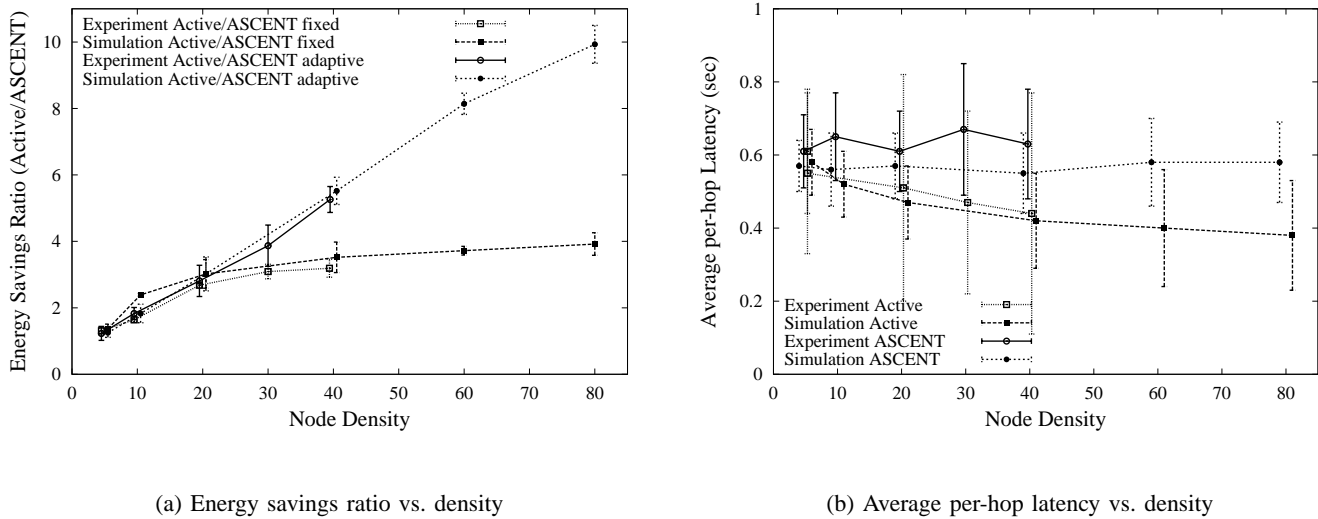


Fig. 9. Figure 9(a) shows the energy savings ratio as a function of density. ASCENT provides significant amount of energy savings over the Active case, up to a factor of 4 with fix timers and 10 with adaptive timers for high density scenarios. Figure 9(b) shows the average per-hop latency as a function of density. ASCENT slightly increases the average hop by hop latency.

Figure 8(a). This is because the end-to-end delivery ratio metric only requires that at least *one* copy of the original packet sent by the source reach the destination. Even in a high-density environment with high losses due to contention for the channel, the likelihood of receiving *one* copy of the packet is still high using flooding.

E. Energy Savings

This section evaluates ASCENT's ability to save energy and increase network lifetime.

In these experiments and simulations, we did not consider the energy spent by the source(s) or the destination(s). For the real experiments, the values are not direct measurements of energy consumption but indirect measurements using the time the nodes spent in the different ASCENT's states.

Figure 9(a) shows the average energy consumption ratio between the active and ASCENT cases as a function of density. We present results using two versions of the ASCENT algorithm, one with fix and the other with adaptive state timers. From these results, we find that ASCENT provides a significant amount of energy savings over the Active case.

When using ASCENT with fixed state timers, we find that as density increases, energy savings do not increase proportionally. This result may seem counterintuitive because in ASCENT the number of active nodes remains constant as density increases, and one would expect to save more energy as the fraction of active nodes decreases. From the analysis shown in Section V-B, we see that the energy consumption, as we increase density, is dominated by the passive-sleep cycle of the passive nodes, and not by the energy consumed by the fraction of active nodes. ASCENT provides a factor of 4 in energy savings in this case.

When using ASCENT with adaptive state timers, we find that as density increases, energy savings do increase proportionally. In this case, nodes can be more aggressive in their sleeping cycle when they detect a high density region, thus increasing the savings they can achieve. Also note that the level of aggressiveness can be tuned based on the probabilistic guarantees offered (number of passive nodes at any given time to react to dynamics).

In both cases (fixed and adaptive), the performance in simulation and real experiments is qualitatively similar but below the expected performance based on the analytical results. The main reason for this is that the analysis done in Section V-B does not consider losses from the environment, which induce ASCENT

to increase the number of nodes with the radio on to maintain a usable topology, and consequently, reduce the energy savings in practice.

F. Latency

We conducted experiments and simulations to evaluate the impact of ASCENT in latency of packet delivery.

Figure 9(b) shows the average per-hop latency as a function of density. Note that we consider only packets that successfully reach the destination in the results (successful end-to-end delivery). We use the average per-hop density to compensate between the different number of hops between the experiments and simulations.

We can see from the graph that ASCENT increases the average per-hop latency when compared to the active case. When using flooding as the routing strategy, the end-to-end delay is affected by the amount of randomization introduced at each hop and the number of nodes forwarding the packets. When density increases, the active case reduces the average per-hop latency because there is a larger probability of a node picking a smaller random interval to forward the packet when there are more forwarding nodes, as it was shown in Section V-B. ASCENT fixes the number of nodes able to forward packets independently of density, and consequently the average per-hop latency tends to remain stable for the same randomization interval.

The reduction in latency for the active case is not as big as predicted in the analytical model. The reason for this is simple: in practice, when considering losses due to the environment and contention for the channel, a packet forwarded fast may not always reach destination, and the average delay per hop can increase from the ideal.

G. Reaction to dynamics

In this section we evaluate how ASCENT react to dynamics introduced by node failures in the active topology.

For these experiments, we let the system run until a stable topology is in place. We then manually kill a set of active nodes such that there is a network partition between the source(s) and destination(s) in the *active* topology.

Figure 10(a) shows the end-to-end delivery rate for ASCENT with both fixed and adaptive state timers. The conditions of the experiment are identical to the experiments performed in Section V-D. The values have been slightly moved on each density point to improve readability. We can see that for the fixed values of α we tested, the end to end delivery rate does not decrease much at high densities. This is because there is high probability that a passive node in the neighborhood exists to fix the *communication hole*. As density decreases, the performance of ASCENT with fixed state timers also decreases. This is because for certain fixed values of α it is possible that we are being over-aggressive in saving energy and all nodes in the neighborhood might be sleeping at the time of the active topology failure. ASCENT with adaptive state timers is more stable for the range of densities we tested.

H. Sensitivity to parameters

This section evaluate the sensitivity of the ASCENT algorithm to the choice of randomization values used in the flooding routing.

Figure 10(b) shows the one-hop delivery rate as a function of density for a larger randomization interval used in flooding. For this experiment we picked a randomization interval of 10 seconds. Figure 8(a) shows a similar graph for a randomization interval of 5 seconds. When comparing the two graphs we can clearly see that for larger randomization intervals we get an increase in the average one-hop delivery rate for different densities. However, there is a trade-off since larger randomization intervals increase the end-to-end latency. For the different levels of randomization we tried, ASCENT case always outperforms the Active case, even when the former has a smaller randomization interval than the latter.

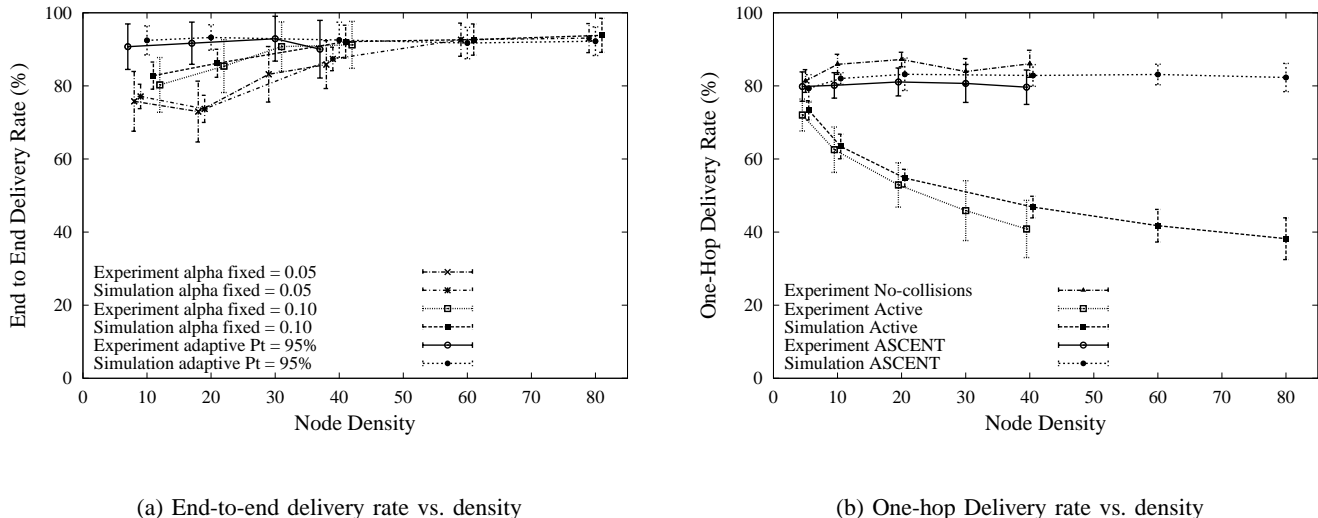


Fig. 10. Figure 10(a) shows the end-to-end delivery rate as a function of density for fixed and adaptive state timers. By using fixed state timers, ASCENT reaction to dynamics performance at lower densities may decrease. When using adaptive state timers ASCENT performance is stable for the range of densities tested. Figure 10(b) shows the one-hop delivery rate as a function of density for a larger flooding randomization interval (10 seconds). ASCENT provides better delivery rates independently of the randomization interval.

The increase is important for the Active case, but it is only marginal for ASCENT. This is because ASCENT operates on a reduced topology independently of the actual density of nodes, and increasing the randomization interval does not help much. This also shows that the expected performance of ASCENT is more stable independently of the choice of the randomization interval.

VI. CONCLUSION AND FUTURE WORK

In this paper, we described the design, implementation, analysis, simulation, 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 packet loss and increase in energy efficiency. Second, ASCENT mechanisms were responsive and stable under systematically varied conditions.

Furthermore, our paper reports on results from experiments using *real* radios, demonstrating the importance of self-configuring techniques that react to the operating conditions *measured locally*.

In the near future, we will evaluate the interactions of ASCENT with new MAC mechanisms, and the use of robust statistical techniques to improve on-line link quality estimation. We will also investigate the use of load balancing techniques to distribute the energy load, and explore the use of wider area links to detect network partitions. More generally, we are interested in understanding the relationships between topology control mechanisms, like ASCENT, and different routing strategies.

This work is an initial foray into the design of self-configuring mechanisms for wireless sensor networks. Our distributed sensing network simulations and 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.

APPENDIX

PROOF OF THE PROBABILITY OF k PASSIVE NODES EQUATION

Given a set of n nodes that alternate between passive and sleep states with probabilities given by equation (6), we would like to find the probability $P(k)$ of at least k passive nodes at any given moment

in time:

$$P(\text{at least } k \text{ passive nodes}) = 1 - P(\text{at most } k - 1 \text{ passive nodes})$$

for $k = 1$:

$$\begin{aligned} P(1) &= 1 - P(0 \text{ nodes passive}) \\ &= 1 - \left(\frac{1}{\alpha + 1}\right)^n \end{aligned}$$

for $k = 2$:

$$\begin{aligned} P(2) &= 1 - (P(0 \text{ nodes passive}) + P(1 \text{ node passive})) \\ &= 1 - \left[\left(\frac{1}{\alpha + 1}\right)^n + \left(\frac{1}{\alpha + 1}\right)^{n-1} \cdot \left(\frac{\alpha}{\alpha + 1}\right) \right] \end{aligned}$$

generalizing for any k :

$$\begin{aligned} P(k) &= 1 - (P(0 \text{ nodes passive}) + P(1 \text{ nodes passive}) + \dots + P(k - 1 \text{ nodes passive})) \\ &= 1 - \left[\left(\frac{1}{\alpha + 1}\right)^n + \left(\frac{1}{\alpha + 1}\right)^{n-1} \cdot \left(\frac{\alpha}{\alpha + 1}\right) + \dots + \left(\frac{1}{\alpha + 1}\right)^{n-k+1} \cdot \left(\frac{\alpha}{\alpha + 1}\right)^{k-1} \right] \\ &= 1 - \left[\left(\frac{1}{\alpha + 1}\right)^n \cdot \alpha^0 + \left(\frac{1}{\alpha + 1}\right)^n \cdot \alpha^1 + \dots + \left(\frac{1}{\alpha + 1}\right)^n \cdot \alpha^{k-1} \right] \\ &= 1 - \left(\frac{1}{\alpha + 1}\right)^n \cdot (\alpha^0 + \alpha^1 + \dots + \alpha^{k-1}) \\ &= 1 - \left(\frac{1}{\alpha + 1}\right)^n \cdot \frac{\alpha^k - 1}{\alpha - 1} \end{aligned}$$

REFERENCES

- [1] K. M. Alzoubi, P.-J. Wan, and O. Frieder. Message-optimal connected-dominating-set construction for routing in mobile ad hoc networks. In *Proceedings of the Third ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, Lausanne, Switzerland, June 9–11 2002. ACM.
- [2] B. Badrinath, Jean Scholtz, Mani Srivastava, Kevin Mills, Vince Stanford, and Eds. Karen Sollins. Special issue on smart spaces and environments. *IEEE Personal Communications*, October 2000.
- [3] J. Broch, D.A. Maltz, D.B. Johnson, Y. Hu, and J. Jetcheva. A performance comparison of multi-hop wireless ad hoc network routing protocols. pages 85–97, October 1998.
- [4] Nirupama Bulusu, John Heidemann, and Deborah Estrin. Adaptive beacon placement. In *Proceedings of the Twenty First International Conference on Distributed Computing Systems (ICDCS-21)*, Phoenix, Arizona, USA, April 16–19 2001. IEEE.
- [5] Gianni Di Caro and Marco Dorigo. Antnet: A mobile agents approach to adaptive routing. Technical Report IRIDIA/97-12, Universite Libre de Bruxelles, Belgium, 1997.
- [6] Alberto Cerpa, Naim Busek, and Deborah Estrin. SCALE: A tool for simple connectivity assessment in lossy environments. Technical Report CENS Technical Report 0021, Center for Embedded Networked Sensing, University of California, Los Angeles (UCLA), September 5 2003.
- [7] Alberto Cerpa, Jeremy Elson, Deborah Estrin, Lewis Girod, Michael Hamilton, and Jerry Zhao. Habitat monitoring: Application driver for wireless communications technology. In *Proceedings of the SIGCOMM Workshop on Communications in Latin America and the Carribean*, San Jose, Costa Rica, April 3–5 2001. ACM.
- [8] Benjie Chen, Kyle Jamieson, H. Balakrishnan, and Robert Morris. Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 85–96, Rome, Italy, July 2001. ACM.
- [9] Jeremy Elson, Solomon Bien, Naim Busek, Vladimir Bychkovskiy, Alberto Cerpa, Deepak Ganesan, Lewis Girod, Ben Greenstein, Thomas Schoellhammer, Thanos Stathopoulos, and Deborah Estrin. EmStar: An environment for developing wireless embedded systems software. Technical Report CENS Technical Report 0009, Center for Embedded Networked Sensing, UCLA, March 2003.
- [10] Deborah Estrin, Ramesh Govindan, John Heidemann, and Satish Kumar. Next century challenges: Scalable coordination in sensor networks. In *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 263–270, Seattle, Washington, USA, August 1999. ACM.
- [11] Deborah Estrin, Ramesh Govindan, and Eds. John Heidemann. Special issue on embedding the internet. *Communications of the ACM*, 43(5), May 2000.
- [12] William Fenner. RFC 2236: Internet group management, version 2, Novemeber 1997.
- [13] Sally Floyd, Van Jacobson, Ching-Gung Liu, Steve McCanne, and Lixia Zhang. A reliable multicast framework for light-weight sessions and application level framing. In *Proceedings of the ACM Special Interest Group on Data Communcation (SIGCOMM)*, pages 342–356, Cambridge, Massachusetts, USA, August 1995. ACM.

- [14] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks. Technical Report UCLA CSD-TR 02-0013, Center for Embedded Networked Sensing, University of California, Los Angeles and Intel Research Lab, University of California, Berkeley, February 2002.
- [15] J. Gao, L. J. Guibas, J. Hershburger, L. Zhang, and A. Zhu. Geometric spanner for routing in mobile networks. In *Proceedings of the 2nd ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 01)*, pages 45–55, Long Beach, California, USA, October 4–5 2001. ACM.
- [16] J. Gao, L. J. Guibas, J. Hershburger, L. Zhang, and A. Zhu. Discrete and computational geometry. *Proceedings of the 2nd ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 01)*, 30(1):45–65, 2003.
- [17] J. L. Gao. *Energy Efficient Routing for Wireless Sensor Networks*. PhD thesis, University of California, Los Angeles, August 2000.
- [18] Jason Hill, Robert Szwedczyk, Alec Woo, Seth Hollar, David Culler, and Kristofer Pister. System architecture directions for networked sensors. In *Proceedings of the Ninth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-IX)*, pages 93–104, Cambridge, MA, USA, November 2000. ACM.
- [19] Chalermek Intanagonwivat, Ramesh Govindan, and Deborah Estrin. Directed diffusion: A scalable and robust communication paradigm for sensor networks. In *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, New York, New York, USA, August 2000. ACM.
- [20] Van Jacobson. Congestion avoidance and control. In *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM)*, pages 314–329, Palo Alto, California, USA, August 1988. ACM.
- [21] David B. Johnson and David A. Maltz. Dynamic source routing in ad hoc wireless networks. *Mobile Computing*, edited by Tomasz Imielinski and Hank Korth, pages 153–181, Kluwer Academic Publishers, 1996.
- [22] Qun Li and Daniela Rus. Sending messages to disconnected users in disconnected ad hoc mobile networks. In *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 44–55, New York, New York, USA, August 2000. ACM.
- [23] Robert Metcalfe and David Boggs. Ethernet: Distributed packet switching for local computer networks. *Communications of the ACM*, pages 395–404, July 1976.
- [24] RFM Monolithics. Tr1000 low power radio system, <http://www.rfm.com/products/data/tr1000.pdf>.
- [25] Vincent Park and Scott Corson. A highly adaptive distributed routing algorithm for mobile wireless networks. In *Proceedings of the Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 1405–1414, Kobe, Japan, April 09–11 1997. IEEE.
- [26] Charles Perkins and Elizabeth Royer. Ad hoc on demand distance vector (aodv) routing. In *Proceedings of the Second IEEE Workshop on Mobile Computing Systems and Applications*, pages 90–100, New Orleans, Louisiana, USA, February 1999. IEEE.
- [27] Gregory Pottie and William Kaiser. Embedding the internet: wireless integrated network sensors. *Communications of the ACM*, 43(5):51–58, May 2000.
- [28] B Prabhakar, E Uysal-Biyikoglu, and A E Gamal. Energy-efficient transmission over a wireless link via lazy packet scheduling. In *Proceedings of the Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 386–394, April 2001.
- [29] Ram Ramanathan and Regina Rosales-Hain. Topology control of multihop wireless networks using transmit power adjustment. In *Proceedings of the Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 404–413, Tel Aviv, Israel, March 26–30 2000.
- [30] Henning Schulzrinne, Stephen Casner, Ron Frederick, and Van Jacobson. RFC 1889: Rtp: A transport protocol for real time applications, January 1996.
- [31] C Schurgers, V Tsiatsis, and M Srivastava. Stem: Topology management for energy efficient sensor networks. In *IEEE Aerospace Conference*, pages 78–89, March 2002.
- [32] Katayoun Sohrabi and Gregory Pottie. Performance of a novel self-organization protocol for wireless ad hoc sensor networks. In *Proceedings of the IEEE Vehicular Technology Conference*, Amsterdam, The Netherlands, September 2000. IEEE.
- [33] Mark Stemm and Randy Katz. Measuring and reducing energy consumption of network interfaces in hand-held devices. *IEICE Transactions on Communications*, E80-B(8):1125–1131, August 1997.
- [34] Yu Wang and Xiang-Yang Li. Geometric spanners for wireless ad hoc networks. In *Proceedings of the 22nd International Conference on Distributed Computing Systems (ICDCS 2002)*, Vienna, Austria, July 2–5 2002. IEEE.
- [35] Alec Woo, Terence Tong, and David Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *Proceedings of the First ACM Sensys Conference*, pages 14–27, Los Angeles, CA, USA, November 5–7 2003. ACM.
- [36] Ya Xu, John Heidemann, and Deborah Estrin. Geography-informed energy conservation for ad hoc routing. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, pages 70–84, Rome, Italy, July 2001. ACM.
- [37] Wei Ye, John Heidemann, and Deborah Estrin. An energy-efficient MAC protocol for wireless sensor networks. In *Proceedings of the Twenty First Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 1567–1576, New York, NY, USA, June 2002. IEEE.
- [38] Yonggang Jerry Zhao and Ramesh Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *Proceedings of the First ACM Sensys Conference*, Los Angeles, CA, USA, November 5–7 2003. ACM.
- [39] Rong Zheng, Jennifer C. Hou, and Lui Sha. Asynchronous wakeup for ad hoc networks. In *ACM International Symposium on Mobile Ad Hoc Networking and Computing*, Annapolis, Maryland, USA, June 2003. ACM.