Topology Control Protocols to Conserve Energy in Wireless Ad Hoc Networks

Ya Xu
USC/Information Sciences Institute
4676 Admiralty Way
Marina Del Rey, CA 90292-6695
yaxu@isi.edu

Solomon Bien
Dept. of Computer Science, UCLA
Los Angeles, CA 90095
sbien@cs.ucla.edu

Yutaka Mori
Dept. of Computer Science, USC
Los Angeles, CA 90089
y.mori@alumni.usc.edu

John Heidemann[1]
USC/Information Sciences Institute
4676 Admiralty Way
Marina Del Rey, CA 90292-6695
johnh@isi.edu

Deborah Estrin
Dept. of Computer Science, UCLA
Los Angeles, CA 90095
destrin@cs.ucla.edu

Alberto Cerpa
Dept. of Computer Science, UCLA
Los Angeles, CA 90095
cerpa@cs.ucla.edu

January 23, 2004

---

[1] Corresponding author

Abstract

In wireless ad hoc networks and sensor networks, energy use is in many cases the most important constraint since it corresponds directly to operational lifetime. This paper presents two topology control protocols that extend the lifetime of dense ad hoc networks while preserving *connectivity*, the ability for nodes to reach each other. Our protocols conserve energy by *identifying redundant nodes* and turning their radios off. Geographic Adaptive Fidelity (GAF) identifies redundant nodes by their physical location and a conservative estimate of radio range. Cluster-based Energy Conservation (CEC) directly observes radio connectivity to determine redundancy and so can be more aggressive at identifying duplication and more robust to radio fading. We evaluate these protocols through analysis, extensive simulations, and experimental results in two wireless testbeds, showing that the protocols are robust to variance in node mobility, radio propagation, node deployment density, and other factors.

Index terms: Wireless sensor networks, Adaptive topology, Topology control, Energy conservation

1 Introduction

Multihop, wireless, ad hoc networking has been the focus of many recent research and development efforts for its applications in military, commercial, and educational environments such as wireless LAN connections in the office, networks of appliances at home, and sensor networks.

A number of routing protocols have been proposed to provide multi-hop communication in wireless, ad hoc networks [27, 6, 28, 26]. Traditionally these protocols are evaluated in terms of packet loss rates, routing message overhead, and route length [7, 24, 13]. A growing emphasis on long-lived networks has added energy consumption as a fourth important metric.

For unattended sensor networks energy consumption is *the* important metric since it maps directly to network operational lifetime. As a preparatory experiment, we compared energy consumption of four common ad hoc routing protocols, (AODV, DSR, DSDV, and TORA), using a simple traffic model in which a few nodes send data over a multi-hop path. This experiment uses an energy [38] and propagation models [7] based on the Lucent WaveLAN direct sequence spread spectrum radio with IEEE 802.11-1997. Full details of this experiment can be found elsewhere [44], but Figure 1 shows the basic results and leads to two observations. First, when idle energy is not considered, reactive protocols (AODV and DSR, the left two dark bars) are much more energy efficient than proactive protocols (TORA and DSDV, the right dark bars), something previously observed. Second and more importantly, when idle-time energy consumption is considered, all protocols are quite similar (the light bars). This second result motivates our need to reduce energy consumed when listening.

In fact, a major source of extraneous energy consumption is from *overhearing* (as previously observed in PAMAS [35]). Radios have a relatively large broadcast range and all nodes in that range must receive each packet to determine if it is to be forwarded or received locally. Although most of these packets are immediately discarded, receiving them consumes energy.

Radios consume power not only when sending and receiving, but also when *listening* or *idle* (the radio electronics must be powered and decoding must occur to detect the presence of an incoming packet). Research [38, 25] shows that idle energy dissipation cannot be ignored in evaluating energy use. Stemm and Katz show idle:receive:transmit ratios of 1:1.05:1.4 by measurement [38], while more recent studies show ratios of 1:2:2.5 [25] and 1:1.2:1.7 [12]. With an energy model that

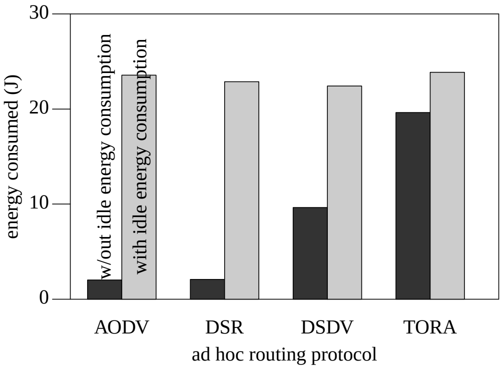

Figure 1: Comparison of energy consumed for four ad hoc routing protocols with different energy models (left, black bars are without considering energy consumed when listening; right, gray bars include this consumption). The simulation has 50 nodes in a 1500m*300m area. Nodes move according to the random way-point model. The energy model is based on Stemm and Katz [38].

takes into account energy use due to overhearing and idle listening, the ad hoc routing protocols we considered consume roughly the same amount of energy (within a few percent) as shown in the gray bars in Figure 1. In the scenario with modest traffic, idle time completely dominates system energy consumption.

The great energy cost associated with idle time and overhearing suggests that energy optimizations must turn off the radio, not simply reduce packet transmission and reception. This approach has been explored at the MAC-level [37, 46, 47, 40]. However, even with MAC protocols that sleep, nodes must behave conservatively, allowing all nodes to listen for possible traffic. These protocols do not take advantage of node density to reduce listen time, and do not exploit application-level information to completely power down nodes.

Our goal is to exploit node redundancy. We observe that when there is significant node redundancy in an ad-hoc network, multiple paths exist between nodes. Thus we can power off some intermediate nodes while still maintaining connectivity. For example, in Figure 2(a), only one of nodes 2, 3, and 4 is required to forward data from 1 to 5: the other two are extraneous. Uninterrupted connectivity between nodes 1 and 5 can be maintained as long as any intermediate node is awake.

This paper describes and analyze two protocols that exploit density to extend lifetime while preserving connectivity. The first, Geographic Adaptive Fidelity (GAF) [45], self-configures redundant nodes into small groups based on their locations and uses localized, distributed algorithms

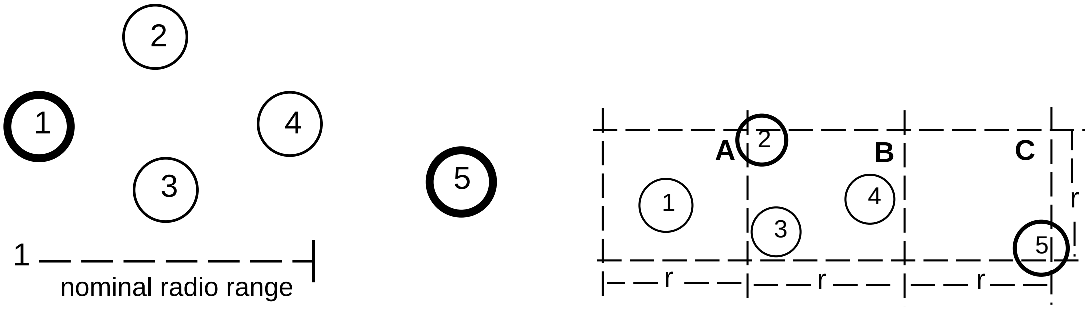

(a) Node Redundancy

(b) GAF Virtual Grid

Figure 2: Examples of node redundancy in ad hoc routing and a GAF virtual grid.

to control node duty cycle to extend network operational lifetime. The second, Cluster-based Energy Conservation (CEC), follows the same principle as GAF, but eliminates its dependency on location information and uniform radio propagation. The challenges of these protocols are how to identify network redundancy, how to control the duty cycle of redundant nodes to conserve energy, and how to maintain connectivity between communicating nodes in a dynamic network when redundant nodes are powered off.

Two additional requirements are needed for successful operation in ad hoc networks. First, the protocol must be *self-configuring*, meaning that it must actively measure the network state in order to react to network dynamics. Second, the protocol must find redundant nodes in a *distributed and localized* fashion, since it is prohibitively expensive to centralize or globally distribute state in rapidly changing ad hoc networks.

In addition to these requirements, a goal is a topology control protocol that is *independent* of the underlying ad hoc routing protocol. Both GAF and CEC can be independent of routing, but experiments show the importance of a routing protocol that can recover from topology changes quickly (see Section 7.1). We also show the relationship between node deployment density and robustness (Section 6.4). Finally, we evaluate GAF and CEC with simple analysis, extensive simulation, and laboratory experiments using real hardware under real radio propagation conditions.

While there is extensive prior work in the area of topology control protocols, most are evaluated by analysis and simulation instead of testbed experiments. The contribution of this paper is therefore the design and development CEC, a protocol that explicitly measures its environment rather than assuming ideal or pessimistic radio propagation or power, and a set of and real-world experiments that demonstrate the importance and efficacy of this approach.

2 Related Work

There has been a significant amount of work in the area of topology control, typically employing analysis or simulation, and considering power control or MAC-level control. We briefly review this work next.

Analytic work: There have been several important analytic 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 [18, 41, 2, 19]. Gao et al. present a randomized algorithm for maintaining a CDS with low overhead [19]. 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 dominated in each grid. The basic algorithm is very similar to GAF, although GAF also considers energy remaining as a parameter for load balancing purposes. Goa's algorithm has an approximation factor of $O(\sqrt{n})$, and 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 [18]. This algorithm also uses the hierarchical clustering algorithm described previously [19]. 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. [2] describe a distributed algorithm for constructing a minimum connected dominating set (MCDS) for the unit-disk-graph with a constant approximation ratio 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 [41]. 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 to define the algorithm constraints. Our work complements theirs by adding the perspective of real-world radio propagation models and evaluation, rather than using only simulation and analysis. Recent work evaluating short-range, low-power radio propagation suggests that noisy and asymmetric connectivity may actually be common or even dominant characteristics of this class of radios [16, 49, 9, 42]. This work suggests there is poor and imprecise correlation between spatial distance and radio connectivity. In fact, connectivity is often not isotropic; "links" are often *asymmetric* and there is not a *monotonic* distance decay of power with distance. Understanding the effects these real-world constraints have on algorithm behavior is important, since a non-monotonic distance/power-relationship or variations in path-loss violate the assumptions in the studies of the above algorithms and may affect correctness. Furthermore, the presence asymmetric connectivity (5-30% of all pairwise communication in [9]) has been observed to cause serious problems with algorithms that assume bidirectional connectivity. Finally, in practice, radio propagation is not circular, thus it is important to validate the above theoretical results against more detailed propagation models, ideally with real-world experiments.

Power control: Reducing energy consumption has been a recent focus of wireless ad hoc network research. One approach has been to adaptively control the transmit power of the radio. LINT/LILT [31] adjusts transmit power in order to form a desired network topology, while the "lazy" scheduling of Prabhakar et al. [30] transmits packets with the lowest possible transmit power for the longest possible time such that delay constraints are still met.

While coarse-grain power control can be very effectively, particularly in asymmetric networks such as cellular telephony, power control may work less well with peer-to-peer, ad hoc or sensor networks [9]. In addition, with short-range radios transmit power represents a small part of the total energy expenditure. With these radios, simply listening can be quite expensive, so we argue that turning the radio off is essential to extending network lifetime. (Future work may explore combinations of these techniques.) Thus, GAF and CEC complement this existing work by exploring this part of the design space. In addition, although GAF and CEC do not require routing protocol information, we show how they can benefit from it, and how they can exploit knowledge of node mobility.

Application-level protocols: By doing energy conservation with *application level* information it is possible to save much more energy, yet the sacrifice is having a network with application-specific characteristics. ASCENT [10] measures local connectivity based on neighbor and packet loss threshold to decide which nodes should join the routing infrastructure based on application requirements, which include the desired level of density and the minimum reception rate in each link. Like CEC, ASCENT is an adaptive measurement-based algorithm that does not depend on geographic assumptions of radio connectivity. PicoNet [4] goes a step farther by designing a system with application-specific hardware and protocols so that energy can be conserved.

MAC-level protocols: Because *MAC level protocols* have a very small view of the network, the main approach followed by such energy-efficient protocols has been to turn off radios that are not actively transmitting or receiving packets. Because it takes generally takes time to turn the radios back on when they are needed, MAC protocols typically trade-off network delay for energy conservation. Energy-efficient MAC and routing protocols can be used together to increase energy conservation.

With TDMA protocols [29], sets of nodes take turns having their radios on and off. PAMAS [35, 36] reduces energy consumption due to overhearing by using a second radio channel to detect activity on its neighboring nodes and turning on its main radio in response to such activity. Sparse Topology and Energy Management (STEM) [34] 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) [46] 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. Newer protocols [**?**] continue to explore MAC-level wake-up schemes.

Routing-level protocols: Like energy-efficient MAC protocols, energy-efficient routing protocols periodically power off nodes; however, such routing protocols do not necessarily cause longer latency. The major difference is that energy-efficient routing *only* powers off the redundant nodes while energy-efficient MAC protocols periodically power off all of the nodes. This difference is due to the fact that unlike MAC protocols, energy-efficient routing protocols have access to network topology information. In addition, routing-level energy conservation protocols can use routing information to ensure that connectivity will be maintained when nodes are turned off.

AFECA [44] seeks to maintain a constant density of active nodes by periodically turning radios off for an amount of time proportional to the measured number of neighbors of a node. By following this approach, as the density increases more energy can be conserved. While AFECA must be conservative in its local density measurement so that network connectivity is not hurt, MIT's SPAN [12] adaptively measures local connectivity in order to makes its decisions. If a node detects that two of its neighbors are not connected by either one or two hops over a network backbone of nodes (by using routing information provided by a routing protocol), then the node joins the backbone itself; otherwise it goes to sleep. The primary difference these protocols and GAF and CEC is the level of integration between the topology control and routing protocols. Both SPAN and AFECA are tightly integrated with the routing protocol, for example, SPAN directly uses routing table state and piggybacking information on routing packets. GAF and CEC instead are independent of the routing protocol and have been tested with several protocols.

There are examples of routing protocols that themselves seek to conserve energy and extend network operational lifetime. Chang and Tassiulas [11], Pottie et al. [29], and LEACH [20] suggest selection of routes based on available energy, so that energy is consumed evenly among nodes and network lifetime is extended. Our approach can complement such efforts.

Finally, our work is related to *adaptive fidelity* [15] and RTCP [33] adaptive frequency techniques. Other examples include beacon density for localization [8] and route fidelity under high mobility [1]. Currently our work maintains a constant fidelity (as do [8, 1]), but future work may explore adaptive fidelity.

3 Geography-informed Energy conservation protocol

With GAF, *Geographic Adaptive Fidelity* [45], nodes that are redundant for communication as determined by geographical position turn off their radios in order to save energy. Nodes alternate having their radios on in order to accomplish load balancing.

3.1 Determining node equivalence

GAF uses location information and an idealized radio model to determine node equivalence. Location information may be provided by GPS or other location systems under development (for example [3, 8, 14]). We have previously evaluated the sensitivity of GAF to errors in location, concluding that GAF is relatively insensitive to moderate random error and large correlated error [45]; we omit those results from here due to space constraints.

Even with location information it is not trivial to find equivalent nodes in an ad hoc network. Nodes that are equivalent for communication between one pair of nodes may not be equivalent for communication between a different pair of nodes. GAF addresses this problem by dividing the whole area where nodes are distributed into small "virtual grids". A virtual grid is defined as follows: for two adjacent virtual grids A and B, all nodes in A can communicate with all nodes in B and vice versa. Thus, in each grid all nodes are equivalent for routing. For example, Figure 2(b) overlays virtual grids on Figure 2(a), creating three virtual grids, A, B, and C. According to our definition of virtual grids, node 1 can reach any of nodes 2, 3, or 4, and nodes 2, 3, and 4 can all reach node 5. Therefore nodes 2, 3, and 4 are equivalent and two of them can go to sleep.

We size our virtual grid based on the nominal radio range R, the farthest possible distance between two nodes in adjacent grids (since they must be able to communicate). If a virtual grid is a square with r units on a side, then the longest possible distance between nodes in adjacent grids is the length of the long diagonal connecting the two grids. Therefore, we get that $r^2 + (2r)^2 \leq R^2$ and thus $r \leq \frac{R}{\sqrt{5}}$. Grid alignment is defined arbitrarily by whatever coordinate system is employed. We make no attempt to optimize grid alignment since with many or moving nodes it seems unlikely that the benefits would outweigh the costs of an informed decision.

3.2 GAF state transitions

In GAF, nodes are in one of three states: *sleeping*, *discovery*, *active*. A state transition diagram is shown in Figure 3. Initially a node starts out in the *discovery* state with its radio turned on and exchanges discovery messages with its neighbors in order to find other nodes within the same grid. The discovery message is a tuple of node ID, grid ID[1], estimated node active time (*enat*),

---

[1] Grid ID is determined by whatever global or relative coordinate system is used.

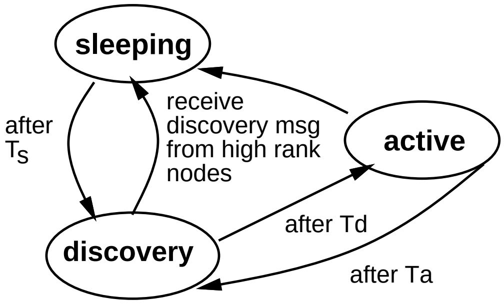

Figure 3: State transitions in GAF.

and node state. If, after waiting for T_d seconds, the node has not already determined that it should not be the active node in its grid, the node moves to the *active* state. A node entering *active* state broadcasts this fact to its neighbors (allowing them to sleep) and then remains active for T_a seconds, then returning to *discovery* state. A node in *discovery* or *active* states can change state to *sleeping* when it can determine some other equivalent node will handle routing. Selection of which node will remain awake is determined by ranking nodes by remaining energy and considering node movement; a node that is active will yield should another node with more energy become active. (See [45] for a more complete discussion of GAF ranking alternatives.) Transitioning to the *sleeping* state, causes a node to cancel all pending timers and power down its radio. A node in the *sleeping* state wakes up after an application-dependent sleep time T_s and transitions back to *discovery*.

For the simulations in this paper we set T_d to be a uniform random value between 0 and a small, fixed time; T_a is adaptively computed as a function of the active node's expected remaining lifetime but limited by its movement speed and capped at 120s; T_s is set to a uniform random value between half and all of the currently active node's advertised T_a. We refer interested readers to [45] for an evaluation of alternative settings for these parameters.

3.3 GAF and Network Mobility

With only a subset of the nodes active, it is possible that network mobility could cause a loss of connectivity. Since GAF nodes know their locations, they can predict when they will leave a grid cell. In GAF-ma (mobility adaptation), each node estimates when it will leave the grid cell and adjusts its *enat* accordingly. More details about GAF-ma can be found in prior work [45].

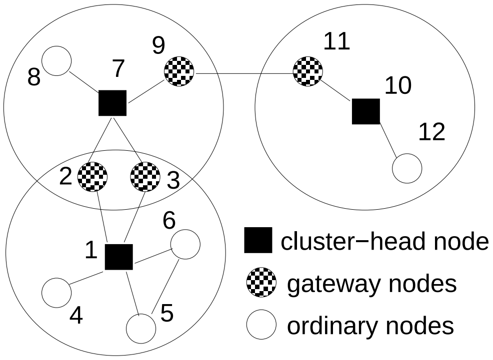

Figure 4: Example of CEC cluster formation. The circle around the cluster-head indicates the radio transmission range. Clusters are interconnected by gateways to provide overall network connectivity.

4 Cluster-based Energy Conservation (CEC) Algorithm

In many settings, such as indoors or under trees where GPS does not work, location information is not available. The dependency on global location information thus limits GAF's usefulness. In addition, geographic proximity does not always lead to network connectivity. As we show in Section 7.1, GAF must make very *conservative* connectivity assumptions because it guesses at connectivity (based on a radio model) instead of directly measuring it. Being conservative requires more nodes to stay active than necessary, leading to less energy conservation.

This motivates *Cluster-based Energy Conservation (CEC)*, which, unlike GAF, does not rely on location information. Further, CEC itself directly and adaptively *measures* network connectivity and thus can find network redundancy more accurately so that more energy can be conserved.

4.1 Determining network redundancy

CEC organizes nodes into overlapping clusters that are interconnected to each other as shown in Figure 4. A *cluster* is defined as a subset of nodes that are mutually "reachable" in at most 2 hops. As shown in Figure 4, a cluster can be viewed as a circle around the cluster-head with the radius equal to the radio transmission range of the cluster-head. Each cluster is identified by one *cluster-head*, a node that can reach all nodes in the cluster in 1 hop.

A *gateway* is a node that connects two clusters, either by being in two clusters (nodes 2 and 3 in Figure 4), a *primary* gateway) or by being adjacent to a gateway in another cluster (nodes 9 and

11 in Figure 4), a *secondary* gateway). The gateway nodes connect all clusters together to ensure overall network connectivity. A node is *ordinary* if it is neither a cluster-head nor a gateway node and is thus redundant.

4.2 Distributed Cluster Formation

In order to elect cluster-heads and gateway nodes, each node periodically broadcasts a *discovery* message that contains its node ID, its cluster ID, and its estimated lifetime. A node's estimated lifetime can be conservatively set by assuming the node will constantly consume energy at a maximum rate until it runs out of energy.

While forming clusters, CEC first elects cluster-heads, then elects gateways to connect clusters. The combination of cluster-heads with gateways guarantees a connected networks as described in Section 4.1.

1. **Cluster-head Selection:** A node selects itself as a cluster-head if it has the longest lifetime of all its neighbor nodes, breaking ties by node ID. Each node can independently make this decision based on exchanged discovery messages. Each node sets its cluster ID to be the node ID of its cluster-head.

2. **Gateway Node Selection:** the gateway nodes, those nodes that can hear multiple cluster-heads are *primary gateway* nodes and those that can hear a combination of cluster-heads and primary gateway nodes are *secondary gateway* nodes.

   When multiple gateway nodes exist between two adjacent clusters, CEC suppresses some of them in order to conserve energy since these gateway nodes are redundant. Gateway selection is determined by several rules. First, primary gateway nodes have higher priority than secondary gateway nodes since at least two secondary gateway nodes, instead of just one primary gateway node, are needed to connect adjacent clusters. Second, gateway nodes with more cluster-head neighbors have higher priority, since this will require fewer nodes to be kept awake. Third, gateway nodes with longer lifetimes have higher priority in order to balance node energy. Note that the gateway selection algorithm does not guarantee that *only* one or one pair of gateway nodes exist between adjacent clusters. In order to support gateway selection, CEC extends the basic discovery message to include the IDs of the clusters that a gateway node can connect.

Figure 4 shows an example of CEC cluster formation in which all nodes have the same estimated network operational lifetime. Nodes 1 and 10 can directly decide they are the cluster-heads because they have the lowest ID of all of their neighbors. Node 7 becomes a cluster-head after nodes 2 and 3 choose node 1 as their cluster-head. Nodes 2 and 3 are primary gateway nodes because they are neighbors of two cluster-heads: nodes 1 and 7. Note that one of nodes 2 and 3 is redundant. Nodes 9 and 11 are secondary gateway nodes between clusters 7 and 10.

4.3 Controlling the Duty Cycle of CEC Nodes

After the selection of cluster-heads and gateway nodes, the remaining redundant nodes are powered off to conserve energy.

Whenever a cluster is formed, each redundant node sets a wake-up timer that will wake it up in time T_s. T_s is set to some fraction of the estimated node lifetime (*enlt*) of the cluster-head. In our CEC implementation, we normally set T_s to be *enlt*/2. In order to avoid thrashing, we set T_s to be *enlt* when it becomes less than a threshold (say 30s). All nodes in the same cluster will thus be powered on to re-form the cluster before the cluster-head runs out of energy, as will gateway nodes in neighboring clusters. While re-forming clusters, it is more likely that the last cluster-head has less remaining energy than the other nodes in the cluster since most have been in a sleeping state and conserving energy. CEC therefore achieves the goal of balanced energy use.

When a node's radio is powered off, its *forwarding* role can be replaced by other nodes. An interesting question is how a sleeping node handles traffic originating from it or destined to it. In the former case, if the node has data to send it can simply power on its radio and send out data. In the latter case, the situation can be addressed as follows. First, for some applications (such as sensor nets), packets are usually not addressed to a particular node, but to a group of nodes with similar properties [22]. Thus, when a node is powered off, other nodes can stay alive to pick up the traffic. Second, some MAC protocols, such as 802.11, support a power-saving mode in which active nodes (typically base-stations) can temporarily buffer data for sleeping nodes. Similar strategies could be used with our schemes. With node interchangeability (as with sensor networks) there should be little penalty, however, if individual nodes must be addressed, evaluating the increase in latency is an area of future work.

4.4 Adapting to Network Mobility

Node mobility can cause network partition. If a cluster-head moves then it might no longer be able to serve as a cluster-head. As with GAF (Section 3.3), CEC uses *mobility prediction* in order to maintain network connectivity.

By estimating how soon a cluster-head will leave its current cluster and informing all nodes in the cluster of that time, the clustered nodes can power themselves on before the cluster-head leaves its cluster. This time is estimated as R/s where s is the cluster-head's current speed and R is its radio transmission range. We assume that nodes can measure their approximate speed, perhaps by wheel odometry, even if they cannot determine their location.

Note that if the R/s estimate is too large, the connectivity between the moving cluster-head and some nodes might be lost before this time. However, if this estimate is too small, CEC will not be able to conserve any energy. In our CEC implementation, we set the estimate as $R/(4s)$ to balance energy conservation and connectivity.

We extend the basic discovery message to include the predicted cluster-leaving time. All nodes in a cluster should wake up to reconfigure clusters before the shorter of T_s and the cluster-leaving time of its current cluster-head. The cluster-leaving time estimate is used analogously in the gateway node selection process (gateway nodes roughly estimate their cluster-leaving time as $R/(2s)$). The additional factors of 4 and 2 make cluster-heads and gateways reevaluate more frequently than regular nodes and are heuristics to reflect their importance.

Although GAF uses a similar method for dealing with mobility, it anticipates hand-offs by using location information, while CEC uses only local measurements. With such global information, GAF may have more accurate mobility predications, but CEC is more practical and localized in nature.

5 Analysis of energy conservation protocols

In order to get an upper bound on how much GAF may extend network lifetime, we consider n nodes that are evenly distributed in an area of size A. The radio of each node has a nominal range of R. As we saw in Section 3.1, the length of a virtual grid square is at most $\frac{R}{\sqrt{5}}$. The minimum number of virtual grid squares, m, would then be $A/(\frac{R}{\sqrt{5}})^2$. Because our nodes are evenly distributed, each grid would have at most n/m nodes, which is equal to $\frac{nR^2}{5A}$ nodes. At best (assuming stationary nodes and no GAF overhead), only one node in each grid will be active while the rest sleep. Based on the maximum number of nodes in each grid, the network lifetime will be extended by at most $\frac{nR^2}{5A}$ times.

The formula reflects the fact that with GAF both more nodes and fewer virtual grids will lead to longer network operational lifetime. The number of virtual grids depends on the nominal radio transmission range and the size of the deployment area.

In order to get the same upper bound for CEC, we again consider n nodes distributed in an area of size A. Again, each node has a radio with a nominal range of R. A cluster area can be viewed as a circle of radius R around the cluster-head. The minimum number of clusters, m, to cover the whole area is equal to $\frac{A}{\pi R^2}$. In each cluster, the cluster-head and a few gateway nodes must stay alive in order to maintain network connectivity. If the average number of adjacent clusters is k, we need at least $m(k+1)$ nodes to cover the whole area. At best (assuming stationary nodes and no CEC overhead), the network lifetime can be extended by $\frac{n}{(m(k+1))}$ times, or $\frac{n\pi R^2}{(k+1)A}$ times.

Since GAF can extend network operational lifetime by at most $\frac{nR^2}{5A}$ times, these equations show that CEC can extend network operational lifetime longer than GAF when $k < 16$. This is a reasonable value for k in most scenarios. One reason for this difference is that GAF conservatively uses smaller grid sizes to group redundant nodes. CEC uses *connectivity measurements* to discover network redundancy and thus does not have the same constraint. A comparison of GAF and CEC is important, however, because GAF is builds on geometric intuition (Figure 2(b)) and thus provides bounds on and insight into CEC performance.

6 Simulation of Topology Control

Because it is difficult to capture the details of GAF/CEC performance in an analytical model, we implemented GAF/CEC in the ns-2.1b6 snapshot of the ns-2 simulator [5] and used AODV and DSR to route packets[2].

We ran GAF/AODV, CEC/AODV, AODV, GAF/DSR, CEC/DSR, and DSR on the same simulated scenarios to compare the effects of variations in node movement, traffic patterns, and energy models on the performance of the protocols, as measured by energy use and data delivery quality.

Traffic, mobility, and radio models: Nodes in the simulation move according to the random way-point model used in [7]. Nodes pause and then move to a randomly chosen location at a fixed speed. We have previously considered [45, 43] pause times ranging from 0 to 900s and node movement rates uniformly distributed between 0 and 20m/s and 0 and 1m/s (representing vehicles and pedestrians, respectively). Due to space constrains, here we present only the extreme cases of 0s and 900s pause times and movement rates between 0 and 20m/s. Nodes move in a 1500m by 300m area.

In most scenarios we use 50 transit nodes that route data (subject to GAF or CEC) and 10 traffic nodes that act as sources and sinks. When we vary node density we use 100 and 200 nodes, while keeping the area constant.

Traffic was generated by continuous bit rate (CBR) sources spreading the traffic randomly among the 10 traffic nodes. The packet sizes were 512 and 1024 bytes and the packet rate was set to four different values: 1 pkt/s, 10 pkts/s, 20 pkts/s and 200 pkts/s. We selected 200 pkts/s to represent an overloaded scenario, since with 1024B packets at that rate, packets alone saturate link bandwidth of 2Mb/s.

We model a radio with a nominal range of 250 meters both with the two-ray-ground propagation model [7] and a non-deterministic shadowing model [32].

---

[2] CMU contributed an extended version of DSR [39] and a validated 802.11 (2M) MAC layer with the simulation package. Our AODV implementation was an improved version from the AODV designers [13]. We have verified that our integration of CMU's ad hoc routing reproduces their results [7], and that our simulation results of unmodified ad hoc protocols are consistent with other published results [7, 13, 24].

Energy model: Our energy consumption model is based on Stemm and Katz's measurements of a 1995 AT&T 2Mb/s WaveLAN (pre-802.11) wireless LAN [38]. They measured costs of 1.6W for transmitting, 1.2W for receiving, and 1.0W for listening. To this we add a cost of 0.025W when sleeping. Newer evaluations of more recent versions of the WaveLAN card and compatible hardware by other vendors show very similar costs [25, 12].

Since it is impossible to evaluate the behavior of the network if the traffic sources and sinks run out of energy before the transit nodes, we give sources and sinks infinite energy. Such nodes follow the same mobility model as transit nodes, but to avoid exploiting their infinite energy sources, they do not run GAF or forward traffic. Because we treat traffic nodes specially, we do not count them when reporting the number of nodes in the simulation.

We give each transit node enough energy so that it can listen for about 450 seconds.

In our GAF simulations we model GPS as consuming 0.033W, the amount of power necessary for reporting location every 8 seconds, since GAF does not require constant position information. We do not turn off GPS when we turn off the radio in order to avoid modeling satellite acquisition time, and because the GPS cost is quite small (about equal to radio sleep cost).

Summary: We conducted our comparison in two phases. In the first phase, we simulated 50 nodes for 900s. Our goal in this phase was to show that our schemes do not reduce the quality of routing, but do in fact conserve energy and extend network operational lifetime. In the second phase we do the same comparison for 3600s while varying the number of nodes in order to see how long network operational lifetime is extended for different node deployment densities.

In each phase we consider 1680 simulations: all combinations of 6 protocols, 7 movement patterns, 10 initial placements, 3 traffic loads, and 2 movement speeds. Based on our results, the difference when running AODV and DSR is not noticeable, so in this section we present results only for the AODV simulations.

6.1 Energy Conservation

In order to quantify energy consumption, we define the *mean energy consumption per node (mecn)* as follows[3]. At the start of the simulation the n nodes have a total initial energy, E_0. After time t,

---

[3] The value of $mecn$ is power; we avoid that term here to avoid confusion with radio transmission power.

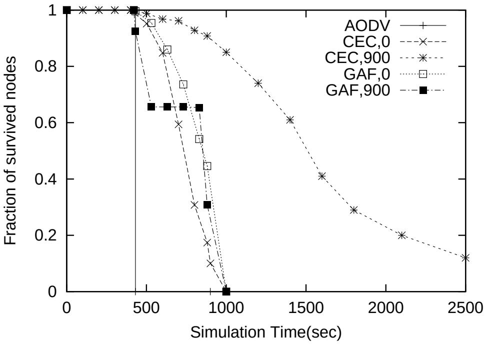

Figure 5: Comparison of non-zero energy node fraction over time: CEC, GAF and plain AODV under different network mobility. Traffic load is 20pkts/s. Different traffic loads do not affect the result. In the legend, "CEC,x" means running CEC with pause time x, so is "GAF,x".

the remaining total energy of the n nodes is E_t. The ($mecn$) equals $\frac{E_0 - E_t}{n * t}$.

Our results show that CEC uses almost half the energy of GAF except in the scenario where nodes move at high speed (20m/s) constantly (zero pause time). GAF typically can save 30-40% more energy than plain AODV, while CEC can save about 60-70% more energy than plain AODV. When nodes move at high speed and with constant movement, CEC adjusts by turning off nodes for shorter times, thus leading to more frequent cluster formations. Such overhead causes CEC to use more energy than GAF in this scenario, though still about 30% less than plain AODV. With the help of global location information, GAF is more energy efficient when dealing with high mobility.

Varying the traffic load does not affect the energy conservation results.

6.2 Extending network operational lifetime

We now examine how these energy savings extend network operational lifetime.

Figure 5 shows the fraction of the network with remaining energy over time when nodes move at 20m/s (When nodes move at low speed, 1m/s, the CEC plot is close to that of the 900s pause time CEC curve, regardless of actual pause time). For both CEC and GAF we plot a zero pause time, representing constant node movement, and a 900s pause time, representing almost no node movement.

All nodes running plain AODV run out of energy at the same time, around 430s. Since AODV does nothing to conserve energy, this result reflects the cost of continuously listening.

We can see that CEC balances energy use more evenly among nodes than GAF. For example,

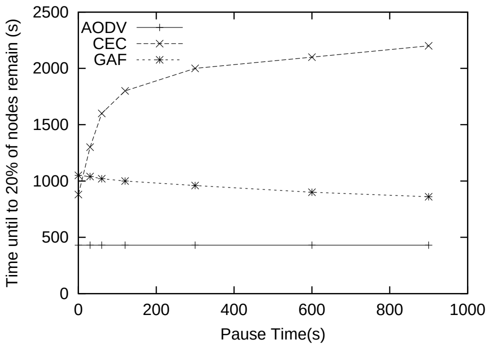

Figure 6: Comparison of network mobility effects on network operational lifetime: CEC, GAF and plain AODV. Traffic load for all scenarios is 20pkts/s.

at time 900s at least 80% of CEC nodes are still alive while at most 40% of GAF nodes are alive except the scenario with a pause time of 0. CEC is more effective at balancing energy because of its connectivity measurement based approach.

CEC also shows a different trend from GAF in regard to mobility. With CEC more nodes survive under low mobility (900s pause time), while in GAF, more nodes survive under high mobility (zero pause time). With CEC, high mobility causes more frequent cluster formations and more overhead. With GAF, high mobility helps balance energy use because changes in node location cause active node re-election within grids. In addition, GAF is more efficient in predicting mobility due to its access to global location information.

In Figure 6 we plot the time at which only 20% of the nodes remain alive against varying degrees of mobility[4]. From this, we can see that CEC extends network operational lifetime at most two times longer than GAF and five times longer than AODV. With CEC, network operational lifetime increases with the pause time. As explained above, this is the effect of the adjustments of CEC for high mobility.

6.3 Network Connectivity

It is easy to conserve energy if one does not care about connectivity: in the extreme, one could turn off the whole network. It is thus important to evaluate the network connectivity produced by our protocols.

---

[4] We chose 20% survival as an arbitrary point of comparision. Selection of too small a value risks skewing the results due to outliers.

We define *data delivery ratio* as the ratio of the number of packets received to the total number of packets sent and the *average data transfer delay* as the the mean delay for those received packets. Under varying traffic loads, these metrics truly reflect the effect of our protocols on network capacity as well as on connectivity.

Within the normal AODV lifetime, we found that CEC performs almost the same as AODV with respect to data delivery ratio, although CEC slightly increases delivery latency at light loads [43]. (We do not show these results here do to space constraints.)

With low network mobility, CEC even performs better than AODV for the following reason. Under heavy traffic, energy use in AODV is unbalanced: those nodes on the routing path or close to the routing path consume more energy and run out of energy sooner than other nodes. The premature loss of these nodes leads to a worse data delivery ratio. This does not happen with CEC since redundant nodes are powered off and node energy use is balanced. This effect does not stand out in high mobility scenarios since different nodes will move in and out of the heavy traffic region–mobility thus leads to balanced energy use.

We also noticed that CEC can maintain almost the same data delivery ratio at extended network operational lifetime. This means that the amount of data carried by the network is doubled.

Figure 7 evaluates data delivery and latencies with GAF and CEC relative to an "ideal", unmodified network. Under high mobility (pause time less than 120s), both CEC and GAF can maintain the same data delivery ratio as the ideal value (with reasonable standard deviation as shown in Figure 7(a)). However, as the pause time increases (larger than 120s), the GAF data delivery ratio becomes worse. At the worst case (pause time 900s), the GAF packet delivery ratio dramatically decreases to only 60% of the ideal level. However, CEC still follows the trend of the ideal data delivery ratio: the lower the mobility, the better the data delivery ratio. The difference between the CEC data delivery ratio and the ideal ratio remains below 5%.

The bad performance of GAF at lower mobility is due to its static gridding mechanism. When node density is high enough to keep at least one node in each grid, GAF works fine. However, when the node density decreases in the extended lifetime, connectivity is affected. With high mobility the situation is not very severe because the movement can help change the uneven distribution.

The same trend is reflected in the delay time as shown in Figure 7(b). CEC follows the trend of the ideal delay time: the lower the mobility, the lower the delay time. GAF performs better than

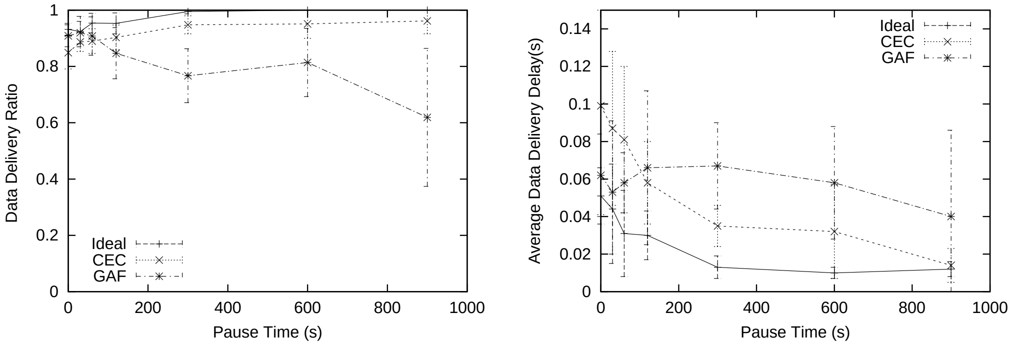

(a) Data Delivery Ratio (b) Average Delay

Figure 7: Data Delivery Quality as a function of pause time comparison: CEC vs. GAF under moving speed 20m/s at extended network operational lifetime. Traffic load is 20pkts/s. Other loads do not change the result. The AODV performance under normal network operational lifetime is plotted as the ideal value.

CEC under high mobility but performs worse at low mobility.

In summary, with a high enough network density, both CEC and GAF can keep good network connectivity. As the network density decreases, CEC can maintain better connectivity than GAF, especially under low mobility.

6.4 Sensitivity to network density

Because our protocols exploit network redundancy, they should extend network operational lifetime farther for a more dense (and more redundant) network. We quantify density as the *number of nodes in nominal radio range (ninra)* [45], so that our results are independent of the number of nodes and the size of the topology.

Figure 8 shows that under high network density (*ninra* larger than 20) both CEC and GAF extend network operational lifetime in proportion to the increase of node density while network operational lifetime under plain AODV remains constant. Under high mobility, CEC and GAF perform about the same; however, under low mobility CEC extends network operational lifetime consistently longer than GAF. With a 4-fold increase in node density (*ninra* 88), CEC extends network operational lifetime 12 times longer than plain AODV and 3 times longer than GAF.

In a low density (*ninra* less than 20) network, a network without redundancy, both CEC and

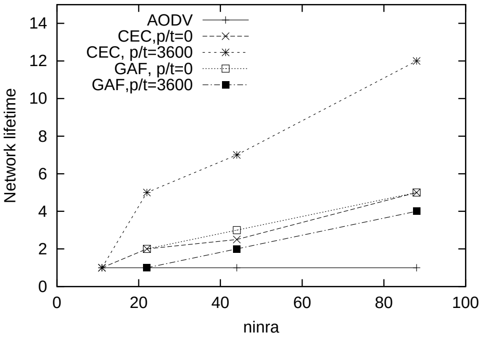

Figure 8: Network operational lifetime comparison among CEC, GAF and plain AODV under different node densities.

GAF have the same network operational lifetime as plain AODV.

When considering the effect of mobility, GAF provides longer network lifetimes at higher mobilities, while CEC does better at lower mobilities. These observations are consistent with our analysis and the ability of GAF to exploit actual locations.

6.5 GAF and CEC Protocol overhead

We measured the energy used by CEC control messages and computed the percentage of energy used by these over the total system energy usage. We find that CEC protocol overhead is always less than 0.4%. Overhead is highest at higher mobility rates. The reason for this is that under high mobility CEC turns on more nodes more frequently. GAF protocol overhead is lower than CEC because less information is required. More details about protocol overhead and mobility can be found in [43].

6.6 Result sensitivity to propagation model

The simulation studies so far have considered a deterministic radio propagation model (the two-ray-ground model). In reality though, radio propagation is strongly affected by multi-path effects (fading). In addition, observations of radio communication in the field show that the shadowing model cannot completely reflect the characteristics of radio propagation. Zhao et al. [51, 52] found that the quality of radio communication between nodes varies dramatically, leading to the belief that there is time-varying interference affecting radio communication in the field.

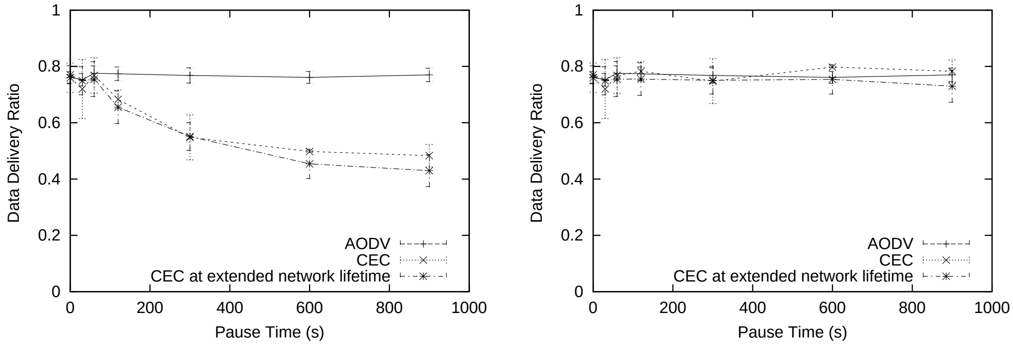

(a) CEC before revision (and AODV) (b) Revised CEC (and AODV)

Figure 9: Packet delivery ratio comparison of CEC and plain AODV under time-varying shadowing model with different pause times. Traffic load is 1pkt/s.

We therefore extend the shadowing model to a time-varying shadowing model by adding a statistical factor to the path loss model so that the attenuation of radio reception changes probabilistically.

We repeated our simulations using our time-varying shadowing model. We chose a value in the range of 3.0 to 4.0 for the path loss exponent in order to reflect a typical outdoor environment and used a standard deviation of shadowing error parameter of 4.0 in the simulation. Our time-varying shadowing model controls how often the path loss exponents should be changed, following an exponential distribution. We observed that 10% to 20% of the links were asymmetric over the simulation time with this model.

As shown in Figure 9(a), under high mobility (pause time less than 120s), AODV, both with and without CEC, has a packet delivery ratio 20% worse with the time-varying propagation model than with the two-ray-ground model. However, the time-varying model does not change the *relative* performance of AODV with and without CEC under high mobility.

Under low mobility (pause time larger than 300s), in contrast, AODV with CEC shows almost a 30% worse packet delivery ratio than that of plain AODV. The reason for the bad performance of CEC under low mobility is that CEC does not sense network topology changes quickly enough. The time-varying shadowing leads to a more frequent change of network topology and has an effect similar to that of high mobility.

CEC performs well under time-varying propagation model when network mobility is high because CEC's mobility prediction algorithm forces CEC to sense the network connectivity more frequently. This frequent measurement ensures that CEC can adapt to any network topology changes, whether a result of mobility or time varying propagation.

The above observations suggest that we can make CEC work more robustly under the time-varying propagation model by having the nodes measure network connectivity more frequently. We can achieve this by adjusting CEC's T_s parameter. Without mobility prediction we set T_s to be *enlt*$/2$, so when a cluster-head had a large amount of energy (a large *enlt*) all nodes in the cluster would check the network connectivity at a very low frequency. To make nodes measure network connectivity more frequently, we can set T_s to be the minimum value of *enlt*$/2$ and a threshold, T_h, decided based on measurement. By doing this, nodes will take measurements at a higher frequency.

We modified our CEC implementation to forcing each node to sense network connectivity on average every 10s (equivalent to a 20m/s node movement speed under high mobility). As shown in Figure 9(b), modified CEC has a statistically equivalent data delivery ratio with this modification. These modifications for more frequent wakeup will reduce CEC's energy conservation, a choice that may be made when nodes have higher mobility.

In summary, the time-varying propagation model introduces the issue of frequent network topology changes even under low mobility. Depending on the application, CEC might need to more frequently measure network connectivity in order to maintain the same data delivery quality as plain AODV.

Similar simulation studies of GAF under harsher propagation models shows that both GAF and AODV show similar reductions in performance as propagation becomes more variable [45].

7 Implementation and Experimentation of Topology Control

While many schemes for topology control have been studied in simulation, it is important to validate these results with real-world experiments. There is growing experimental evidence that short-range, low-power radios can have very non-uniform radio propagation. Packet-level simulations must necessarily use simplified radio propagation models, perhaps resulting in optimistic conclusions. Through experimentation with GAF, we demonstrate below that protocols such that require assumptions about radio propagation require very conservative assumptions to work in the field.

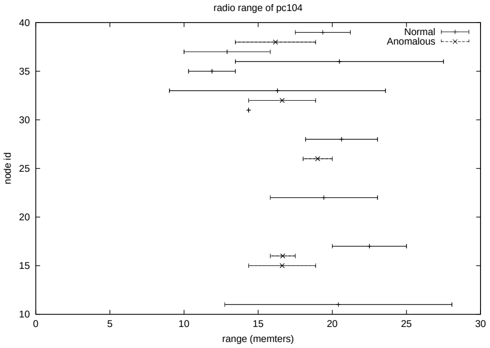

Figure 10: Variance of radio range.

These results suggest that protocols like CEC that adapt based on pairwise node connectivity can perform better in practice.

7.1 GAF with Experimental Radios

GAF's off-line construction of virtual grids based on radio range is a cause for concern, since it assumes a circular radio propagation model, in which the radius of the circle is the range of the radio. In order to test the validity of making such use of radio ranges, we did an experiment in which we measured the connectivity between pairs of nodes. Based on this measure and some chosen thresholds, pairs of nodes are classified as well-connected (95% successful packet reception between the two nodes) or poorly-connected (10% successful reception). When comparing the connectivity measure to the actual distance between a pair of nodes, we see very high variability. Figure 10 shows an approximation of the radio range of each node. (Recent work has studied this phenomena in detail [50, 42].) For each node, the left side of the bar shows the shortest distance to a poorly-connected node and the right side of the bar shows the longest distance to well-connected node. In order to trust radio ranges, we would want these two distances to have close to the same value. However, for only a few nodes (nodes 15, 16, 26, 32, and 38) is this actually the case. Such variance in radio range is also observed in [17]. Because radio ranges cannot be accurately estimated by distance, GAF is difficult to configure–a conservative setting of R results in grids too small to provide energy conservation.

We evaluated GAF in a network of 15 PC/104s equipped with Radiometrix packet radios, running directed diffusion as the routing protocol [22]. We considered three configurations: small network size with fixed R, larger network with fixed R, and manually chosen non-uniform R. Given our observations about variation in radio range, we chose a conservative R for the fixed scenarios. As a result, the larger network with fixed R had several empty grid cells, though the network could still function. Due to space constraints we do not present detailed results here, but in general energy savings was proportional to the number of nodes per grid cell (densities were about 1.5 for fixed R and 2.7 for manual configuration). Thus configuration with non-uniform grids provided better energy savings. While GAF was effective, the difficulty at selecting good values for R suggests that CEC can provide better performance in practical settings.

We also observed an interaction between GAF and the routing protocol. Since they operate independently, in some cases GAF would turn off a node that was actively routing packets. Thus in some cases we observed periods of interrupted communication and very high latency. These problems were not observed in simulations of AODV and DSR because those routing protocols include a local repair mechanism for broken routes. We are currently adding such a mechanism to directed diffusion.

7.2 CEC Implementation and Experimental Setup

By *measuring* connectivity instead of assuming the correlation of radio range to distance, CEC provides a solution to the problem of variable radio ranges.

We ran experiments with CEC on a testbed of 21 iPAQs, each equipped with a UCB mote [21] as a radio interface and an 802.11 card for experimental control and logging purposes. The nodes were arranged in a square grid, with nodes at most of the vertices of the grid. Each node has an average of seven neighbors.

Data packets are routed through the use of flooding: each node keeps a cache of the data packets it has received and, with each subsequent arrival of a packet, forwards (broadcasts) that packet only if it is not already present in the cache. All nodes in the network generate traffic according to the following rule: every five seconds, each node probabilistically decides whether to generate a new data packet. The probabilities used generate an expected one new packet per five seconds over the whole network.

The energy usage of each node is modeled in the same way as in our simulations. Again, we only model the energy usage of the radio, since the energy usage of other components is assumed to be negligible in comparison. Each node is given enough energy to remain in idle listening mode for 450 seconds.

Our hope is that these experiments will provide qualitative confirmation of our simulation results. We could not directly control many parameters: we lacked an AODV implementation for our experimental platform, so instead we chose flooding because its simplicity avoids complications due to the specifics of the routing protocol. We did not have enough real hardware to test scenarios as large as our simulations. And of course we could not carefully control radio propagation. However, as described below, we believe the results confirm qualitatively our simulations over real hardware.

7.2.1 CEC Extension of Network Operational Lifetime

Figure 11(a) depicts the extension of network lifetime by showing the number of nodes with remaining energy over time. Without CEC, there is a sharp drop at time 450s to zero, when all nodes run out of energy. With CEC, though, we see that network operational lifetime (20% of nodes remaining) is extended until time 1000s. The curve is qualitatively similar to what we observed in Figure 5, validating those results. Further, the fact that the curve is smooth indicates that CEC successfully balances energy usage among equivalent nodes.

7.2.2 CEC Data Delivery Ratio

As we pointed out in the discussion of the simulations, it is important to look both at energy savings and quality of data delivery, since optimizing for one hurts the other. There are two elements to this metric. First, we must judge the connectivity of the active nodes. Second, we must see how well the active nodes cover the network.

Figure 11(b) shows the data delivery ratio among CEC active nodes over time as well as the percentage of the network that is active. The first thing to notice is that the plot of the number of active nodes receiving packets and the plot of the number of active nodes are about the same. This means that about 100% of the active nodes are receiving the data packets and that the set of active nodes is indeed connected. The second thing to notice is the fluctuations in the bottom two plots. The first reason for these is the fact that nodes in discovery phase are counted as active:

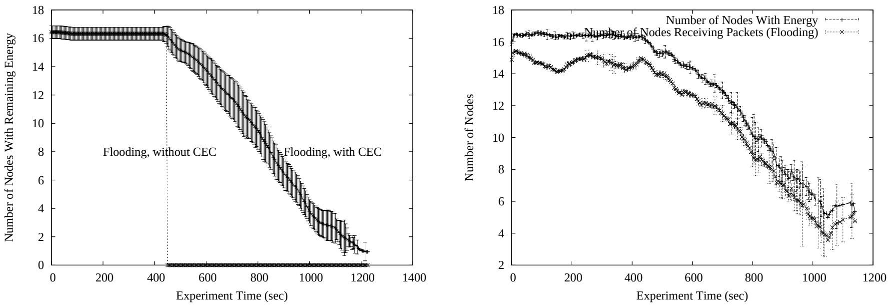

(a) Extension of network operational lifetime (95% confidence interval)

(b) CEC active nodes over time (95% confidence interval)

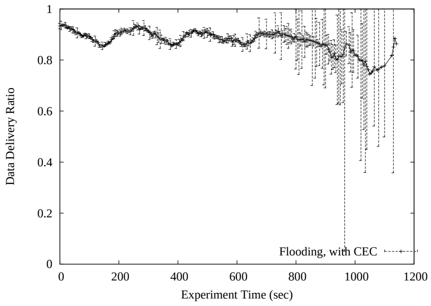

(c) Data delivery ratio (95% confidence interval)

Figure 11: CEC experimental results.

during discovery phase (e.g. just after time 200s) a larger number of nodes will temporarily be active. While the periodic discovery phases amplify the fluctuations in this plot, the fluctuations are present even if one does not count nodes in discovery state as active. CEC chooses active nodes based on two factors: a node's remaining energy and a node's number of neighboring nodes. In the first iteration of CEC, all nodes have approximately equal amounts of energy. As a result, CEC chooses active nodes primarily based on the number of neighboring nodes of a node. In the first iteration, therefore, the number of active nodes is close to the smallest possible for the given network topology. However, in the second iteration these same nodes (which make up something close to the minimum dominating set of nodes) cannot be picked (since they have used up half of their energy in the first iteration). Thus, a different and larger set of active nodes must be chosen in the second iteration.

Figure 11(c) shows the second part of the data delivery metric, the data delivery ratio as a function of time. We compute data delivery as number of currently reachable nodes that successfully receive flooded data. This value would be less than 1 because of packet loss due to collisions or corruption that occur in flooding to currently awake nodes, or due to network partition. We assume that awake nodes will buffer packets to their sleeping neighbors. We observe that delivery ratio is generally high. Although CEC turns off a subset of nodes, data is still effectively delivered throughout the network.

Two interesting features of the plot are the dips in the number of nodes receiving packets around times 200s and 400s. By looking at Figure 11(b), one can see that the dips occur during discovery phases. We attribute these dips to the fact that all nodes are taking part in the flooding of data packets. There is thus a higher probability of collisions with all of the nodes forwarding packets than there is with only the active nodes forwarding packets.

Starting just before time 800s, we see increased variance in the data plotted. This is to be expected as the nodes in the network run out of energy, since the network at this point might be left disconnected. In a disconnected network, a packet could get to vastly different amounts of the network depending on the topological location of the node that generates the packet.

Something that we noticed is that packets often reached either every node in the network or none of the nodes in the network. In Figure 11(c) we have ignored the packets that reach none of the network. Such packets are not even successfully transmitted over the first link. We suspect that perhaps by lowering the packet generation rate, we can eliminate such packets.

In our experiments with real radios we have seen that CEC is an effective scheme for saving energy without sacrificing the usefulness of networks for delivering data. CEC eliminates the redundancy that can be safely given up.

8 Future Work

There are several areas for future work.

We qualitatively compared our protocols with other similar protocols, but it was not possible to do a quantitative comparison because of differences in protocol development environments and lack of implementations of similar protocols. Future work quantitatively comparing GAF and CEC with SPAN and similar protocols would be helpful.

In addition, it is important to evaluate how these “layer 2.5” topology control protocols interact with other sleep/wakeup protocols, particularly at the MAC-level. Our expectation is that MAC-level sleep/wakeup would reduce the impact of our schemes, but that the combination would show greater energy savings, particularly at very dense deployments.

Our evaluation considers only a very simple mobility model (random waypoint). Future work should consider richer models and larger scenarios, particularly in light of recently improved mobility models [23, 48].

9 Conclusions

We have described two new protocols, GAF [45] and CEC. Both control node topology to extend overall network lifetime, GAF via geographic information, CEC by directly measuring connectivity. Through simulation we demonstrated that in stationary networks, GAF can extend 2- to 4-fold and CEC 6- to 12-fold at medium or high network densities (40 or 90 neighbors). We described extensions to both GAF and CEC improve reliability in the face of high mobility and found 2- to 4-fold lifetime extension for both protocols there.

Furthermore, our paper reports on what are (to our knowledge) the first experimental results for this class of topology control protocols. These experiments demonstrate that radio propagation in today's sensor networks is too variable to make GAF feasible for deployment, demonstrating the importance of adaptive observation of network connectivity.

References

[1] S. Ahn and A. U. Snakar. Adapting to route-demand and mobility (ARM) in ad-hoc network routing. In *9th International Conference on Network Protocols*, pages 56–66, Nov. 2001.

[2] K. M. Alzoubi, P.-J. Wan, and O. Frieder. Message-optimal connected-dominating-set construction for routing in mobile ad hoc networks. In *Third ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2002)*, Lausanne, Switzerland, June 9–11 2002. ACM.

[3] P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *Proceedings of the IEEE Infocom*, pages 775–784, Tel-Aviv, Israel, March 2000.

[4] F. Bennett, D. Clarke, J. B. Evans, A. Hopper, A. Jones, and D. Leask. Piconet: Embedded mobile networking. *IEEE Personal Communications Magazine*, 4(5):8–15, Oct. 1997.

[5] L. Breslau, D. Estrin, K. Fall, S. Floyd, J. Heidemann, A. Helmy, P. Huang, S. McCanne, K. Varadhan, Y. Xu, and H. Yu. Advances in network simulation. *IEEE Computer*, 33(5):59–67, May 2000. Expanded version available as USC TR 99-702b at `http://www.isi.edu/~johnh/PAPERS/Bajaj99a.html`.

[6] J. Broch, D. B. Johnson, and D. A. Maltz. The dynamic source routing protocol for mobile ad hoc networks. *INTERNET-DRAFT, draft-ietf-manet-dsr-03.txt*,, October 1999. Work in progress.

[7] J. Broch, D. Maltz, D. Johnson, Y. Hu, and J. Jetcheva. A performance comparison of multi-hop wireless ad hoc network routing protocols. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 85–97, October, 1998.

[8] N. Bulusu, J. Heidemann, and D. Estrin. GPS-less low cost outdoor localization for very small devices. *IEEE Personal Communications Magazine*, 7(5):28–34, Oct. 2000.

[9] A. Cerpa, N. Busek, and D. 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.

[10] A. Cerpa and D. Estrin. ASCENT: Adaptive self-configuring sensor network topologies. In *Twenty First International Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, June 2002.

[11] J.-H. Chang and L. Tassiulas. Energy conserving routing in wireless ad-hoc networking. In *Proceedings of the IEEE Infocom*, pages 22–31, Tel Aviv, Israel, Mar. 2000. ACM/IEEE.

[12] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris. Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. *ACM Wireless Networks*, 8(5), September 2002.

[13] S. R. Das, C. E. Perkins, and E. M. Royer. Performance comparison of two on-demand routing protocols for ad hoc networks. In *Proceedings of the IEEE Infocom*, pages 3–12, Tel Aviv, Israel, March 2000.

[14] L. Doherty, K. S. J. Pister, and L. E. Ghaoui. Convex position estimation in wireless sensor networks. In *Proceedings of the IEEE Infocom*, pages 1655–1663, Alaska, April 2001. IEEE.

[15] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Next century challenges: Scalable coordination in sensor networks. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 263–270, Seattle, Washington, USA, Aug. 1999. ACM.

[16] 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.

[17] 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, UCLA Computer Science, Jul., 2002.

[18] 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.

[19] 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.

[20] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy-efficient communication protocols for wireless microsensor networks. In *Hawaii International Conference on System Sciences*, pages 2–12, Jan. 2000.

[21] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister. System architecture directions for networked sensors. In *Proceedings of the 9th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS IX)*, pages 93–104, Cambridge, MA, Nov, 2000.

[22] C. Intanagonwiwat, R. Govindan, and D. Estrin. Directed diffusion: A scalable and robust communication paradigm for sensor networks. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 70–84, Boston, Massachusetts, USA, Aug, 2000.

[23] A. Jardosh, E. M. Belding-Royer, K. C. Almeroth, and S. Suri. Towards realistic mobility models for mobile ad hoc networks. In *Proceedings of the ACM International Conference on Mobile Computing and Networking*, pages 217–229, San Diego, California, USA, Sept. 2003. ACM.

[24] P. Johansson, T. Larsson, N. Hedman, B. Mielczarek, and M. Degermark. Scenario-based performance analysis of routing protocols for mobile ad-hoc networks. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 195–206, August, 1999.

[25] O. Kasten. Energy consumption. ETH-Zurich, Swiss Federal Institute of Technology. Available at `http://www.inf.ethz.ch/˜kasten/research/bathtub/energy_consumption.html%`, 2001.

[26] V. Park and S. Corson. Temporally-ordered routing algorithm (TORA) version 1 functional specification. *INTERNET-DRAFT, draft-ietf-manet-tora-spec-02.txt,*, October 1999. Work in progress.

[27] C. Perkins. Ad hoc on demand distance vector (AODV) routing. *Internet-Draft, draft-ietf-manet-aodv-04.txt*, pages 3–12, October 1999,Work in progress.

[28] C. E. Perkins and P. Bhagwat. Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers. In *Proceedings of the ACM SIGCOMM*, pages 234–244, Auguest 1994. A revised version of the paper is available at `http://www.cs.umd.edu/projects/mcml/papers/Sigcomm93.ps`.

[29] G. J. Pottie and W. J. Kaiser. Embedding the internet: wireless integrated network sensors. *Communications of the ACM*, 43(5):51–58, May 2000.

[30] B. Prabhakar, E. Uysal-Biyikoglu, and A. E. Gamal. Energy-efficient transmission over a wireless link via lazy packet scheduling. In *Proceedings of the IEEE Infocom*, pages 386–394, Apr. 2001.

[31] R. Ramanathan and R. Rosales-Hain. Topology control of multihop wireless networks using transmit power adjustment. In *Proceedings of the IEEE Infocom*, pages 404–413, March 2000.

[32] T. S. Rappaport. Wireless communications: Principles and practice. *Prentice Hall, Upper Sadde River, New Jersey 07458*, Reprinted 1999.

[33] H. Schulzrinne, S. Casner, R. Rrederick, and V. Jacobson. RTP: A transport protocol for real-time application. RFC 1889, IETF, Jan. 1996.

[34] C. Schurgers, V. Tsiatsis, and M. Srivastava. STEM: Topology management for energy efficient sensor networks. In *IEEE Aerospace Conference*, pages 78–89, March, 2002.

[35] S. Singh and C. Raghavendra. PAMAS: Power aware multi-access protocol with signalling for ad hoc networks. *ACM Computer Communication Review*, 28(3):5–26, July 1998.

[36] S. Singh, M. Woo, and C. Raghavendra. Power-aware routing in mobile ad hoc networks. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 181–190, October, 1998.

[37] K. Sohrabi and G. J. Pottie. Performance of a novel self-organization protocol for wireless ad hoc sensor networks. In *Proceedings of the IEEE Vehicular Technology Conference*, page xxx, Amsterdam, The Netherlands, Sept. 2000. IEEE.

[38] M. Stemm and R. H. Katz. Measuring and reducing energy consumption of network interfaces in hand-held devices. *IEICE Transactions on Communications*, E80-B(8):1125–1131, Aug. 1997.

[39] The CMU Monarch Project. The CMU monarch projects wireless and mobility extension to ns. `http://www.monarch.cs.edu`.

[40] T. van Dam and K. Langendoen. An adaptive energy-efficient mac protocol for wireless sensor networks. In *Proceedings of the First ACM SenSys Conference*, pages 171–180, Los Angeles, California, USA, Nov. 2003. ACM.

[41] Y. Wang and X.-Y. 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.

[42] A. Woo, T. Tong, and D. 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, California, USA, Nov. 2003. ACM.

[43] Y. Xu. *Adaptive Energy Conservation Protocols for Wireless Ad Hoc Routing*. PhD thesis, University of Southern California, 2002.

[44] Y. Xu, J. Heidemann, and D. Estrin. Adaptive energy-conserving routing for multihop ad hoc networks. Technical Report TR-2000-527, USC/Information Sciences Institute, Oct. 2000. Available at `ftp://ftp.isi.edu/isi-pubs/tr-527.pdf`.

[45] Y. Xu, J. Heidemann, and D. Estrin. Geography-informed energy conservation for ad hoc routing. In *Proceedings of the ACM International Conference on Mobile Computing and Networking*, pages 70–84, Rome, Italy, July 2001. ACM.

[46] W. Ye, J. Heidemann, and D. Estrin. An energy-efficient MAC protocol for wireless sensor networks. In *Proceedings of the IEEE Infocom*, pages 1567–1576, New York, NY, USA, June 2002. USC/Information Sciences Institute, IEEE.

[47] W. Ye, J. Heidemann, and D. Estrin. Medium access control with coordinated, adaptive sleeping for wireless sensor networks. *ACM/IEEE Transactions on Networking*, 2003. accepted to appear IEEE/ACM Transactions on Networking; draft available as ISI-TR-567.

[48] J. Yoon, M. Liu, and B. Noble. Sound mobility models. In *Proceedings of the ACM International Conference on Mobile Computing and Networking*, San Diego, California, USA, Sept. 2003. ACM.

[49] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *Proceedings of the First ACM SenSys Conference*, pages 1–13, Los Angeles, California, USA, Nov. 2003. ACM.

[50] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *Proceedings of the First ACM SenSys Conference*, pages 1–13, Los Angeles, California, USA, Nov. 2003. ACM.

[51] Y. Zhao, R. Govindan, and D. Estrin. Aggregate network properties for monitoring wireless sensor networks, 2002. Unpublished Manuscript.

[52] Y. Zhao, R. Govindan, and D. Estrin. Residual energy scans for monitoring wireless sensor networks. In *Proceedings of IEEE Wireless Communications and Networking Conference (WCNC'02)*, pages 78–89, Florida, USA, March, 2002.