Geography-informed Energy Conservation for Ad Hoc Routing *

Ya Xu
Information Science Institute
4676 Admiralty Way, Ste 1001
Marina del Rey, CA 90292
yaxu@isi.edu

John Heidemann
Information Science Institute
4676 Admiralty Way, Ste 1001
Marina del Rey, CA 90292
johnh@isi.edu

Deborah Estrin
Comp. Sci. Dept., University of California, Los Angeles
3713 Boelter Hall
Los Angeles, CA 90095
estrin@cs.ucla.edu

ABSTRACT

We introduce a *geographical adaptive fidelity* (GAF) algorithm that reduces energy consumption in ad hoc wireless networks. GAF conserves energy by identifying nodes that are equivalent from a routing perspective and then turning off unnecessary nodes, keeping a constant level of routing *fidelity*. GAF moderates this policy using application- and system-level information; nodes that source or sink data remain on and intermediate nodes monitor and balance energy use. GAF is independent of the underlying ad hoc routing protocol; we simulate GAF over unmodified AODV and DSR. Analysis and simulation studies of GAF show that it can consume 40% to 60% less energy than an unmodified ad hoc routing protocol. Moreover, simulations of GAF suggest that network lifetime increases proportionally to node density; in one example, a four-fold increase in node density leads to network lifetime increase for 3 to 6 times (depending on the mobility pattern). More generally, GAF is an example of *adaptive fidelity*, a technique proposed for extending the lifetime of self-configuring systems by exploiting redundancy to conserve energy while maintaining application fidelity.

1. INTRODUCTION

Multihop, ad hoc networking has been the focus of many recent research and development efforts. Wireless networks and multihop routing have application in military, commercial, and educational environments including wireless office LAN connections, home networks of devices, and sensor networks.

A number of routing protocols have been proposed to provide multi-hop communication in wireless, ad hoc networks [22, 4, 23, 21]. Traditionally these protocol are evaluated in terms of packet loss rates, routing message overhead, and route length [5, 15, 10]. Since ad hoc networks will often be deployed using battery-powered nodes, comparison and optimization of protocol energy consumption is also important (as suggested for future work by some researchers [15]).

When ad hoc networks are deployed using battery-powered nodes, the important question of how limited energy resources affects system lifetime and overall performance becomes critical. For scenarios such as sensor networks where energy use maps directly to lifetime and utility, energy use is *the* important metric. To understand energy efficiency we examined existing ad hoc routing protocols using models of Lucent WaveLAN direct sequence spread spectrum radio with the IEEE 802.11-1997 protocol with representative models of energy consumption [30] and radio propagation [5]. We first only consider energy cost due to packet transmission or reception. Such costs may also include energy dissipation in MAC-level retransmissions, RTS/CTS etc. We studied energy consumption of four ad hoc routing protocols (AODV, DSR, DSDV, and TORA) with a simple traffic model where a few nodes send data over a multihop path [32] (Figure 1). With this energy model we found that on-demand protocols such as AODV and DSR consume much less energy than a priori protocols such as DSDV (the left, dark bars in Figure 1). A priori protocols are constantly expending energy pre-computing routes, even though there is no traffic passing on these routes.

In other words, on-demand protocols, by their very nature, are more efficient in the energy consumed by routing overhead packets. As a result, energy use is dominated by routing protocol overhead. In fact, the major source of extraneous energy consumption was from *overhearing*, as previously observed in PAMAS [28]. Radios have a relatively large broadcast range. 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, they consume energy with this simple energy model. This observation motivates approaches that avoid overhearing. The PAMAS protocol suggests a MAC-layer approach to minimize this cost [29]; TDMA protocols would also be applicable (for example [24]).

Actual Radios consume power not only when sending and

*This research is supported by the Defense Advanced Research Projects Agency (DARPA) through the SCADDS project under DARPA contract DABT63-99-1-0011.

To appear in the Proceedings of the Seventh Annual ACM/IEEE International Conference on Mobile Computing and Networking(ACM Mobicom), July 16-21, 2001, Rome, Italy.

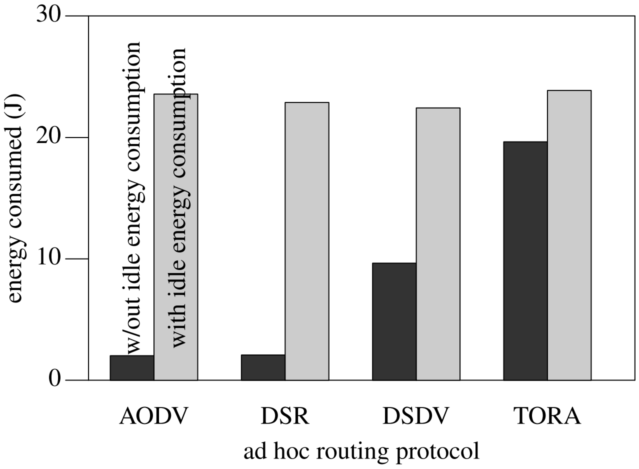

Figure 1: Comparison of energy consumed for four ad hoc routing protocols with different energy models (left, black bars are without considering energy consumed when listening; right, gray bars include this consumption). The simulation has 50 nodes in a 1500m*300m area. Nodes move according to the random way-point model. The energy model is based on Stemm and Katz [30].

receiving, but also when *listening* or *idle* (the radio electronics must be powered and decoding to detect the presence of an incoming packet). Research [30, 16] shows that idle energy dissipation can not be ignored in comparing to sending and receiving energy dissipation. Stemm and Katz show idle:receive:transmit ratios are 1:1.05:1.4 by measurement [30], while more recent studies show ratios of 1:2:2.5 [16] and 1:1.2:1.7 [9]. In any of these cases, energy dissipation in idle state can not be ignored. With such energy model, all ad hoc routing protocols considered consume roughly the same amount of energy (within a few percent) as shown in the grey bars on Figure 1. In the scenario with modest traffic, idle time completely dominates system energy consumption.

The studies based on an energy model that considers energy dissipation in sent/received packets and idle time, suggest that energy optimizations must turn off the radio, not simply reduce packet transmission and reception. Powering off radio conserves energy both in overhearing due to data transfer, and in idle state energy dissipation when no traffic exists. We therefore explore nodes that power down their radios much of the time. This approach is similar to the use of TDMA for power conservation [24], or PAMAS [28]. However, unlike these approaches, we employ information from above the MAC-layer to control radio power. (We make use of the power management controls in IEEE 802.11 to control power[1].) The application- or routing-layers provide better information about when the radio is not needed.

[1]802.11 supports power saving mode in both infrastructure network and ad hoc network. Note that powering on/off MAC is just like node moving in/out communication range with other nodes. RTS/CTS is still used in unicast communication to address hidden terminal issue. [30] shows that time for 802.11 MAC on/off is in a few milliseconds. In other words, powering on/off MAC does not affect normal 802.11 operation.

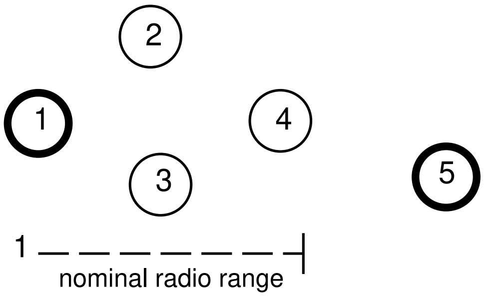

Figure 2: Example of node redundancy in ad hoc routing.

On the other hand, we observed 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, if node 2 is awake, nodes 3 and 4 are extraneous for communication between 1 and 5. We define *routing fidelity* as uninterrupted connectivity between communicating nodes. Thus routing fidelity (that 1 and 5 can communicate) can be maintained as long as any intermediate node is awake.

These observations motivate the two primary contributions of our design:

1. The use of application- and system-information to turn off node radios for extended periods time. Node duty cycles are influenced by application endpoints and node movement patters to preserve communication fidelity. Predictions about node lifetimes allow energy-conscious load balancing.

2. The use of *node deployment density* to adaptively adjust routing fidelity. Routing redundancy is correlated with denser node deployment (when many nodes can hear each other). We show how to use this information to increase node duty cycles and to extend the lifetime of the network as a whole.

   We employ *location information* (such as from the Global Position System, GPS) and active node communication to determine node density and redundancy. (We use GPS or other localization methods to determine density; we are also exploring density determination through communications alone as future work.)

The GAF approach is one example of *adaptive fidelity* [12] and RTCP [27] adaptive frequency techniques. GAF keeps the fidelity of network reachability constant while adapting node behavior to extend network lifetime. Other examples include beacon density for localization [6]. More generally, we wish to design self-configuring networks that exploit redundancy to conserve energy while preserving the fidelity of network applications.

2. ENERGY-CONSERVING ROUTING ALGORITHM

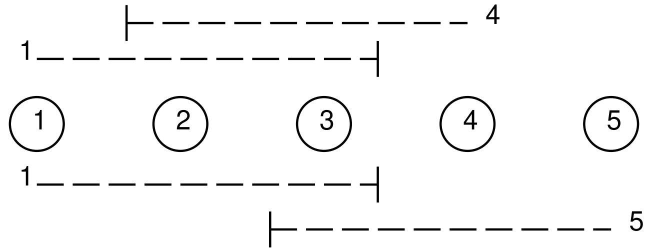

Figure 3: The problem of node equivalence.

We next present our energy-conserving ad hoc routing algorithm, GAF, *Geographical Adaptive Fidelity*. Each GAF node uses location information to associate itself with a "virtual grid", where all nodes in a particular grid square are equivalent with respect to forwarding packets. Nodes in the same grid then coordinate with each other to determine who will sleep and how long (Section 2.2); this determination is moderated by application and system information (Section 2.3 and 2.5). Nodes then periodically wake up and trade places to accomplish load balancing (Section 2.4). We also consider how GAF interacts with the underlying ad hoc routing protocol in Section 2.6.

2.1 Determining node equivalence

GAF uses location information and virtual grids to determine node equivalence. Location information used in GAF may be provided by GPS or other location systems under development (for example [1, 6, 11]). For our initial discussion, we assume that each node knows its current location exactly relative to other nodes. In Section 3.9 we relax this assumption and show that GAF is not affected by moderate location error or even by large, correlated error.

Even with location information, it is not trivial to find equivalent nodes in an ad hoc network. Nodes that are "equivalent" between some nodes may not be equivalent for communication between others. For example, in Figure 3 nodes are equidistant 1 unit apart with radio range slightly larger than 2 units. For communication between nodes 1 and 4, nodes 2 and 3 are equivalent, while between 1 and 5 only node 3 is acceptable.

GAF addresses this problem by dividing the whole area where nodes are distributed into small "virtual grids". The virtual grid is defined such that, for two adjacent grids A and B, all nodes in A can communicate with all nodes in B and vice versa. Thus all nodes in each grid are equivalent for routing. In GAF, nodes exchange grid IDs to adjust their duty cycle. For example, Figure 4 overlays virtual grids on Figure 2, creating three virtual grids A, B, and C. According to our definition of virtual grids, node 1 can reach any of 2, 3, or 4, and 2, 3, and 4 can all reach 5. Therefore nodes 2, 3, and 4 are equivalent and two of them can sleep.

In the definition of virtual grid, we require that any node in adjacent grid can *communicate* with each other. In reality, a node's radio communication range is not deterministic or even symmetric due to radio propagation effects such as multi-path reflection. In our initial discussion, we assume that the communication range is deterministic (using the tworayground propagation model frequently used in many ad hoc routing studies, for example [5, 15]). In Section 3.8, we compare the effects of non-deterministic radio propagation using a shadowing model, finding that shadowing propagation models do not change our comparisons between GAF and ad hoc routing protocols.

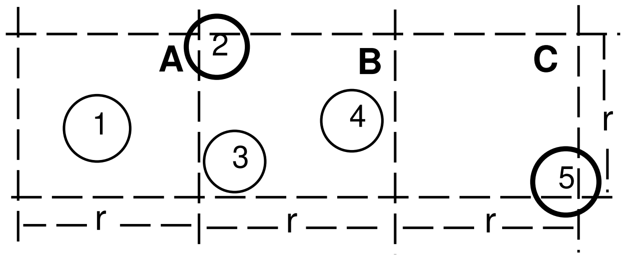

Figure 4: Example of virtual grid in GAF.

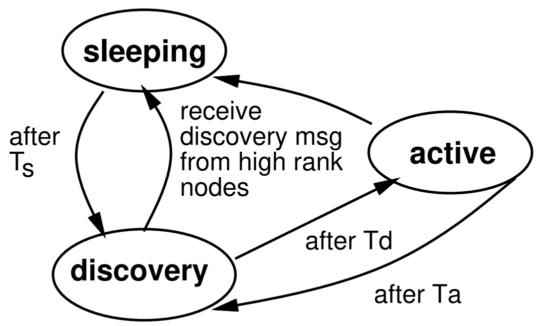

Figure 5: State transitions in GAF.

We size our virtual grid based on the nominal radio range R. Assume virtual grid is a square with r units on a side as shown in Figure 4. In order to meet the definition of virtual grid, the distance between two possible farthest nodes in any two adjacent grids, such as grid B and C in Figure 4, must not be larger than R. For example, node 2 of grid B and node 5 of grid C in Figure 4 are at the end of the long diagonal connecting two adjacent grids. Therefore, we get:

$$r^2 + (2r)^2 \leq R^2 \tag{1}$$

or

$$r \leq \frac{R}{\sqrt{5}} \tag{2}$$

2.2 GAF state transitions

In GAF, nodes are in one of three states: *sleeping*, *discovery*, *active*. A state transition diagram is shown in Figure 5.

Initially nodes start out in the *discovery* state. When in state *discovery*, a node turns on its radio and exchanges discovery messages to find other nodes within the same grid. The discovery message is a tuple of node id, grid id, estimated node active time (*enat*), and node state. As described above, a node uses its location and grid size to determine the grid id.

When a node enters *discovery* state, it sets a timer for T_d seconds. When the timer fires, the node broadcasts its discovery message and enters state *active*. The timer can also be suppressed by other discovery messages. This timer reduces the probability of discovery message collision.

When a node enters *active*, it sets a timeout value T_a to define how long this node can stay in *active* state. After T_a, the node will return to *discovery* state. While active, the node periodically re-broadcasts its discovery message at intervals T_d.

A node in *discovery* or *active* states can change state to *sleeping* when it can determine some other equivalent node will handle routing. Nodes negotiate which node will handle routing through an application-dependent ranking procedure described in the next section. (Node ranking can be an arbitrary ordering of nodes to decide which nodes should be active, or it can be selected to optimize overall system lifetime.) When transitioning to *sleeping*, a node cancels all pending timers and powers down its radio.

A node in the *sleeping* state wakes up after an application-dependent sleep time T_s and transitions back to *discovery*.

2.3 Tuning GAF

GAF leaves choices of many parameters including *enat*, T_d, T_a, node rank, T_s to applications. In this section, we describe how and why these parameters are chosen in the current GAF algorithm. Applications may wish to optimize these choices, for example, perhaps trading increased packet loss for greater energy savings.

Estimated node active time (enat) can be set to the expected node lifetime (*enlt*), conservatively set by assuming the node will constantly consume energy at a maximum rate until it dies. Rather than this conservative *enat*, GAF uses an approach described in Section 2.4 to balance energy usage across nodes.

GAF selects the *discovery message interval* (T_d) as a uniform random value between 0 and some constant. This approach avoids contention from synchronized discovery messages (as inspired by SRM [13]). The range of T_d can also be influenced by node rank to encourage highly ranked nodes to suppress low-ranked nodes, allowing them to rapidly go to sleep. Nodes in the *active* state may wish to chose a larger T_d to avoid bandwidth and energy overhead.

Nodes *active duration* (T_a) can be its expected lifetime (*enlt*). GAF instead uses T_a to accomplish load balancing as described in Section 2.4.

Node ranking in GAF is chosen to maximize network lifetime by selecting which nodes handle routing. Rank is determined by several rules. First, A node in the *active* state has higher rank than a node in *discovery* state. This rule tries to quickly reach the state each grid only maintains one active node. For nodes with the same state, GAF gives nodes with longer expected lifetime (*enat*) higher rank. This rule put nodes with longer expected lifetime into use first. Finally, node ids are used to break ties. It is possible to let applications choose different node rank rules according to their own bias, for example, application may favor the active nodes until they drain out energy in a sensor network.

Node *sleep duration* (T_s) can be set to the *enat* of the active node since this is the conservative assumption of its lifetime. Due to node mobility, the *active* node may move out the grid (of course there is chance that other nodes move into this grid). This can leave a grid without any *active* nodes although some nodes are sleeping, reducing routing fidelity. One approach that statistically reduces this problem is to set T_s as a uniform random time between 0 and *enat*. This large range of T_s may often have nodes wake up quite early. In GAF therefore T_s is uniformly from the range $[enat/2, enat\ ]$. We also consider an alternate approach using node mobility information in Section 2.5.

2.4 Load balancing energy usage

GAF employs a load balancing strategy so that all nodes remain up and running together for as long as possible. The idea behind this is that all nodes in the network are equally important and no one node must be penalized more than any of the others. (An alternative is to completely exhaust the energy of each node in turn while other nodes sleep.)

GAF uses the following load balancing strategy. After a node remains in the *active* state for time Ta, it changes its state to *discovery* to give a chance to other nodes within the same grid to become *active* (figure 5). Recall that nodes are ranked according to their remaining energy levels. When the active node changes its state to *discovery*, it is more likely that it has less remaining energy than its neighbor nodes because presumably the neighbors were in the *sleeping* state conserving energy during the node's active time. Consequently, the node that was active is less likely to remain active after the discovery phase.

The active node sets Ta to the value *enat* and advertises *enat* in its discovery messages. The non-active nodes in the neighborhood use *enat* to determine their sleeping period. The active node sets *enat* to a value less than the time to use up all remaining energy (*enlt*). In our simulations we set *enat* to *enlt*/2 so that the node consumes half of its energy before handing off to another node in the neighborhood.

To avoid thrashing, when *enlt* becomes less than a threshold (say, 30s) GAF sets *enat* to the full *enlt*.

2.5 Adapting to high mobility

GAF tries to adapt the number of nodes participating in ad hoc routing to keep a constant level of nodes that route data. The ideal scenario would be one active node in each grid at any time. However, as nodes move, the active node may leave its grid. This may leave the prior grid without an active node, reducing routing fidelity. In scenarios with high mobility this problem can greatly increase packet drop rates.

We can accommodate high mobility by considering this system-level behavior explicitly in GAF. Each node estimates the time it expects to leave its grid (the *expected node grid time* or *engt*) and includes this information in the discovery message. When other node enter *sleeping* state, they sleep for the smaller of *enat* and *engt* to decide how long it can stay in *sleeping* state. This change does not change the node rank; nodes use the same ranking rules to decide *who* sleeps, but they sleep for shorter time.

A node can estimate *engt* based on its current speed s (speed can be obtained or estimated from most GPS receivers) and the grid size, $engt = r/s$. This estimate works well for the mobility model we use (the random way-point model with pauses); in systems where movement is less predictable this value may be more difficult to estimate.

In order to compare the effect of the node mobility adaption, we call the GAF without node mobility adaption as GAF-basic (GAF-b), the GAF with node mobility adaption as GAF-mobility adaption (GAF-ma) and use simulation to compare their performance. GAF-b and GAF-ma behave the same when mobility is low. When mobility is high, GAF-b will tend to have fewer active nodes and therefore lower energy consumption but higher packet loss rates as shown in Section 3.5.

2.6 GAF interactions with ad hoc routing

In principle, GAF will run over any ad hoc routing protocols because it only uses application- and system-level information to decide each node's duty cycle, and since discovery messages are only broadcast to direct neighbors. In later sections we evaluate GAF combined with AODV [22] and DSR [5] against unaugmented AODV and DSR. For brevity, we will sometimes say "we compare GAF and AODV" in place of "we compare AODV with GAF with unmodified AODV".

GAF decision to turn nodes on and off is independent of ad hoc routing protocols. If a node is actively routing packets when it is powered off, GAF depends on the ad hoc routing protocol quickly re-routing traffic. This may cause some packet loss, although most of ad hoc routing protocols react to changes quickly. An optimization that we have not explored is to have GAF inform the ad hoc routing protocol of impending suspension, allowing it to preemptively re-route any traffic.

3. PERFORMANCE EVALUATION

To evaluate our schemes, we first use a simple mathematical analysis to determine an idealized level of energy conservation in GAF. Since analysis cannot capture the complexity in a full GAF scenario, we then use simulation to study GAF effects on network lifetime, how and why it conserves energy, and whether or not it increases the number of packet drops. Finally, we show that network lifetime under GAF is proportional to the density of node deployment, and we examine the sensitivity of our simulations due to error in the radio propagation model and quality of location.

3.1 Analytic performance analysis

To get an upper bound on how much GAF may extend network lifetime we next consider a very simple analytic model. Assume that n nodes are evenly distributed in a area with topography size A. Nominal radio range for each node is R. According to Equation 2, the grid size can be set as $\frac{R}{\sqrt{5}}$ which is the maximum size of a virtual grid. The minimum total number of virtual grid cells, m, would be

$$m = \frac{A}{(\frac{R}{\sqrt{5}})^2} \tag{3}$$

According to our assumption of evenly distributed nodes, each grid would have at most n/m nodes, or

$$\frac{n * R^2}{5 * A} \tag{4}$$

nodes.

At best, assuming stationary nodes and no GAF overhead, only one node in each grid will be active while the rest sleep. Since Equation 4 gives the maximum number in each grid, the network lifetime will be extended at most $(n*R^2)/(5*A)$ times.

The formula basically reflects the fact that with GAF algorithm, the more nodes, the longer network lifetime, and the fewer number of virtual grids, the longer network lifetime. The number of virtual grids mainly depends on the nominal radio transmission range and the topography size. Equation 2 gives the upper-bound of grid size.

The overhead due to GAF discovery message is small. Although GAF periodically sends out discovery message if the node is in the *discovery* or *active* state, the frequency will be very low. Since the broadcast is limited in one hop around a node, such overhead will not affect the whole system energy dissipation too much.

In the following sections we use simulation to relax our assumptions and explore GAF performance in more realistic conditions.

3.2 Simulation methodology

It is difficult to capture the details of GAF performance in an analytical model. For that reason we implemented GAF in the ns-2 simulator [3], evaluating variations in node movement, traffic pattern and energy model as described below.

In order to demonstrate the flexibility of GAF, we implemented GAF in a snapshot of ns-2.1b6. We use locally modified and extended version of ad hoc routing contributed by CMU [31][2], an improved AODV implementation from the AODV designers [10], and a energy model described below. We attached GAF to AODV to get GAF/AODV, and GAF to DSR to get GAF/DSR. We then run GAF/AODV, AODV, GAF/DSR, DSR on the same simulated scenarios to compare the performance in terms of energy dissipation and data delivery quality. We have verified that our integration of the CMU's ad hoc routing reproduces their results [5], and that our simulation results of unmodified ad hoc protocols are consistent with other published results [5, 10, 15].

[2] CMU also provided a validated 802.11 (2M) MAC layer simulation with the simulation package

Traffic and mobility models: Nodes in the simulation move according to the random way-point model used in CMU [5]. Nodes alternate between pausing and then move to a randomly chosen location at a fixed speed. We consider seven pause times: 0, 30, 60, 120, 300, 600, and 900 seconds. For each pause time we generate 10 sets of initial placements and random way-points. We also evaluate two different node movement speeds: uniform distribution between 0 and 20m/s and uniform distribution between 0 and 1m/s. Nodes move in a 1500m by 300m meter area.

In most scenario we use 50 transit nodes (nodes running ad hoc routing) and 10 traffic nodes acting as sources and sinks. When we examine node density, we vary the number of nodes from 50 to 100 and 200 while keeping area constant.

Simulation traffic was generated by continuous bit rate (CBR) sources spreading the traffic randomly among 10 traffic nodes. The packet size was 512 bytes. The packet rate was set to three different values: 1 pkt/s, 10 pkts/s and 20 pkts/s to evaluate GAF sensitivity to traffic load. Even with 20 pkts/s packet rate, the traffic load is still well below the network capacity. We chose to study light to moderately loaded systems because nodes in ad hoc networks are expected to be energy constrained, more so than bandwidth constrained.

We model a radio with a nominal range of 250 meters with both the two-ray-ground propagation model and a shadowing model.

Energy model: Our energy consumption model is based on Stemm and Katz's measurements of an 1995 AT&T 2Mb/s WaveLAN (pre-802.11) wireless LAN [30]. They measured costs of 1.6W for transmit, 1.2W for receiving, and 1.0W for listening. To this we add a cost of 0.025W when sleeping. We chose their model as representative at the time we began this work. Although this hardware is now somewhat old, newer evaluations of more recent versions of the WaveLAN card and compatible hardware by other vendors shows very similar costs. Specifically, Kasten measured the Digitan 2Mb/s 802.11 wireless LAN, observing costs of 1.9W (transmit), 1.5W (receive), 0.75W (listen), 0.025W (sleep) in 2000 [16], and Chen et al. measured the Lucent 2Mb/s WaveLAN 802.11 cards, observing costs of 1.4W (transmit), 1.0W (receive), 0.83W (listen), and 0.043W (sleep) [9]. In all of these studies of 802.11 hardware, transmission is 50–100% more expensive than listening, while sleeping is a tiny fraction of this cost.

It is impossible to evaluate the behavior of the network as whole if traffic nodes run out energy before the transit nodes. To avoid this artifact we separate the traffic nodes from routing nodes and give traffic nodes infinite energy. Traffic nodes follows the same mobility model as transit nodes, but they do not run GAF, nor do they forward traffic. Because we treat traffic nodes specially, we do not count them when reporting the number of nodes in the simulation (for example, our 50 node simulation consists of 60 nodes, 50 transit nodes and 10 traffic nodes). Note that our scheme does not exclude the possibility of accommodating limited-energy traffic nodes as well. In this case, node in the *active* state would need to buffer data for sleeping end nodes.

We give each transit node enough energy so that it can listen for about 450 seconds. Since nodes in AODV and DSR have nodes listening constantly, all nodes expire after 450s even without traffic.

Since we depend on GPS, we must model GPS energy consumption. We model GPS as consuming 0.46W in continuous reporting mode, with power conservation modes that consume 0.165W by reporting every second and 0.033W reporting every 8 seconds [16]. Since GAF does not require constant position information, we add a GPS cost of 0.033W to GAF simulations (and not to simple AODV or DSR simulations). We do not turn off GPS when we turn off radio to avoid modeling satellite acquisition time, and because GPS cost is quite small (about equal to radio sleep cost).

Summary: We compared GAF/AODV with normal AODV and GAF/DSR with normal DSR. We conducted our comparison into two phases. In the first phase, we simulated 50 nodes for 900s (twice of AODV/DSR lifetime) to see how GAF affects network lifetime, data delivery, and energy usage. Our goal in this phase is to show that GAF does not reduce the quality of ad hoc routing and that it will conserve energy compared to an unmodified ad hoc routing protocol. We then repeat the same comparison, varying the number of nodes (50, 100, and 200), for 3600s of simulation time to observe what will happen when all nodes run out of energy. In this second phase, we study how long GAF will extend network lifetime.

In each phase, We consider 1680 simulations, all combinations of 4 protocols (GAF/AODV, AODV, GAF/DSR, DSR), 7 movement patterns, 10 initial placements, 3 traffic loads, and 2 movement speeds (1m/s and 20m/s).

The remainder of this section presents our simulation results. Sections 3.3–3.5 analyze GAF performance in terms of network lifetime, energy conservation, and data delivery quality under low mobility with various traffic loads. Sections 3.6–3.9 investigate GAF sensitivity to various scenario parameters, namely high mobility (3.6), node density (3.7), shadowing propagation model (3.8), and location error (3.9).

3.3 GAF extension of network lifetime

The first question we examine is how GAF affects network lifetime. We measure network lifetime as the fraction of transit nodes with non-zero energy as a function of time. In this section we fix the number of nodes at 50 and consider the range of pause times and the two mobility speeds. We present the results for AODV with traffic loads at 20 pkts/s. We also evaluated DSR and other traffic loads; we do not present those results as they were similar. (Traffic load does not substantially affect network lifetime because the energy used in packet forwarding is much less than the cost of idle-time listening.)

In Figure 6 we present the results for AODV and GAF-b with movement of 1m/s and traffic loads at 20 pkts/s. GAF-b and -ma are equivalent at this movement speed. We leave the discussion of high mobility effect to section 3.6.

The first thing we observe in Figure 6 is that with AODV all

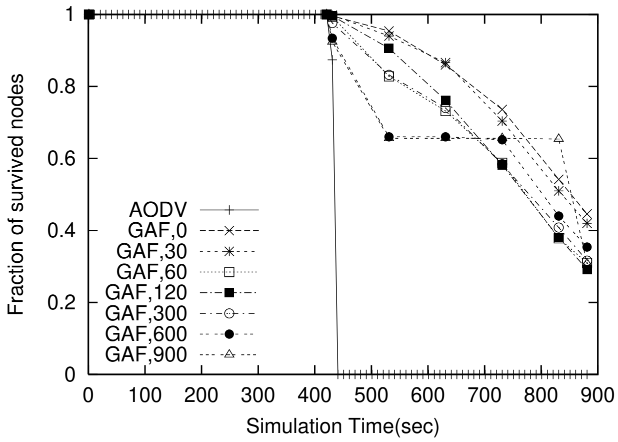

Figure 6: Network lifetime comparison: GAF vs. AODV at low node speed (1m/s) under various pause time. Movement: 1m/s, traffic: 20 pkts/s.

nodes run out of energy at the same time at about 450s. This is a function of the "battery" in each node and the fact that AODV does nothing to conserve energy. This represents the cost of continuously listening.

GAF extends network lifetime considerably; after 900s 30–40% of nodes are still alive depending on mobility pattern. Scenarios with shorter pause times consistently have better network lifetime than those with longer pause time. This is because moving nodes are better at load balancing. Consider the extreme case of a 900s pause time (no movement during the simulation). In this case, grid cells with a single node must remain constantly active and so will expire at 450s. Grids with two nodes will last longer and then all die at the same time producing the stair-steps at 550s and 850s. When nodes move (with pause times less than 600s), they will sometimes move into these grids and share the load, thus producing more gradual node failure.

3.4 GAF energy savings

Node lifetime is a useful measure, but it can be a crude measure of actual energy consumption because node life is binary, so 50 about-to-die nodes are considered as good as 50 fresh nodes. In order to quantify how much energy GAF saves, we instead compute the mean energy consumption per node.

To define the mean energy consumption per node, we assume that simulation starts with n nodes with initial total energy of n nodes as E_0, After time t, the remaining total energy of n nodes is E_t. Then the mean energy consumption per node (aen) is:

$$aen = \frac{E_0 - E_t}{n * t} \tag{5}$$

We calculate aen for both AODV, GAF-b and GAF-ma for all different scenarios including node moving speed, movement pattern and different traffic load.

As we did in section 3.3, we only discuss the low mobility scenario here. We leave the discussion of high mobility into section 3.6.

The calculation of aen shows that at low node move speed (1m/s), both GAF-b and GAF-ma have about 40% lower aen than AODV no matter what movement patterns (different pause time) are. The different traffic load does not affect the average energy consumption per node too much. The reason is although traffic load is increasing, over the long simulation time, idle energy dissipation still dominates the total energy dissipation in the whole system.

3.5 GAF effects on data delivery

While we have shown that GAF conserves energy, because GAF causes nodes to sleep it may reduce how many packets are successfully delivered, reducing routing fidelity. We have designed GAF to maintain a constant level of routing fidelity, but we expect some data loss when nodes go to sleep as routes change, and we have identified several cases where node dynamics may cause unintentional data loss.

We define two metrics to measure data delivery. The first metric is the data delivery ratio, which is the ratio of the number of received packets over the total sent packets. The second metric is the average data transfer delay, which the the mean delay for those received packets.

In this section, we compare GAF and AODV during the first 450s of simulation when all AODV nodes are alive. After this time AODV fails to deliver any packets; considering the entire time would therefore skew these measures in GAF's favor. Our goal here is to show that GAF is not significantly worse than AODV when AODV is effective. In the section 3.7 we will evaluate these metrics over the whole lifetime of GAF.

We calculate the data delivery ratio and average delay time for both AODV, GAF-b and GAF-ma for all different scenarios including node moving speed, movement pattern and different traffic load.

At low node speed (1m/s), both GAF and AODV have the same data delivery ratio of 99% and the same mean delay across all pause times. These results are because at low speeds, nodes rarely moves out of their virtual grids.

We leave the discussion on high mobility to section 3.6.

3.6 Effect of high mobility on GAF-b and GAF-ma

We summarize GAF-b and GAF-ma performance in terms of network lifetime extension, energy saving, and data delivery quality under high mobility in this section.

Network lifetime extension We evaluated GAF with a higher movement speed (20m/s). With more movement we expect to see differences in GAF-b and -ma performance. In the scenarios of high node moving speed, GAF-b and GAF-ma should have different behavior in network lifetime because GAF-ma uses high mobility adaption to keep conservative on powering off nodes while GAF-b aggressively powers off nodes. We therefore plot network lifetime of

GAF-b and GAF-ma in high node moving speed in two different Figures: Figure 7(a) is for GAF-b and Figure 7(b) is for GAF-ma.

Figure 7(a) and Figure 7(b) share the same characteristics as we discussed in Figure 6. First, AODV lifetime is not affected in all scenarios; secondly, GAF extends the network lifetime in different degrees: the less pause time, the longer network lifetime; and thirdly, traffic load does not change the network lifetime too much.

The difference between Figure 7(a) and Figure 7(b) is that GAF-b can keep more nodes alive for a longer time, especially in the high mobility pattern. For example in Figure 7(a), there are still 90% of nodes are still alive at the time 900s when the pause time is 0. There is only about 40% survived in GAF-ma in the same scenario (See Figure 7(b)). The reason, as we discussed before, is that GAF-b aggressively powers node off for energy conservation while GAF-ma conservatively powers off the nodes in high mobility.

GAF energy savings At high node move speed (20m/s), GAF-b and GAF-ma have the same 40% lower *aen* than AODV at long pause time (more than 300 seconds); GAF-b and GAF-ma has incrementally reduced *aen* along with shorter pause time. When pause time is 0, GAF-b has 60% lower *aen* than AODV, and GAF-ma has 50% lower *aen* than AODV. GAF-b shows its greater energy conservation on the high mobility scenario which conforms to our findings as we just described above.

Data delivery quality At high node moving speed (20m/s) we see some differences between GAF-b and GAF-ma (see Figure 8). At long pause times (more than 300 seconds) GAF-b, GAF-ma, and AODV continue to have the same data delivery ratio and mean delay. At shorter pause time (less than 300 seconds) GAF-ma and AODV remain similar, but GAF-b has worse data delivery ratio and longer average delay. At the most extreme case with a 0 pause time, GAF-b has an 85% data delivery ratio, while GAF-ma and AODV can reach 95%. The average delay for GAF-b is 0.35 second. when both GAF-ma and AODV are around 0.16 second.

GAF-b's node sleep policy is the reason for its worse behavior under high or constant mobility. We address this problem three ways: first, GAF-b may be unsuitable at high mobility under the current node density; Second, by considering mobility in the sleep policy, GAF-ma is more appropriate than GAF-b in high mobility scenarios. Third, so far we have only considered comparison within the AODV lifetime. Since typically GAF has much longer lifetime than AODV, over this extended lifetime GAF can perform much better than AODV. We consider extended lifetimes in the next section. Finally, we show later (Figure 9(a)) that at higher node density GAF-b is more robust to packet loss.

3.7 How network density affects GAF

GAF extends network lifetime by identifying equivalent routing nodes and putting these equivalent routing into different duty cycles. As node density increases we expect GAF to take advantage of the additional equivalent nodes to extend

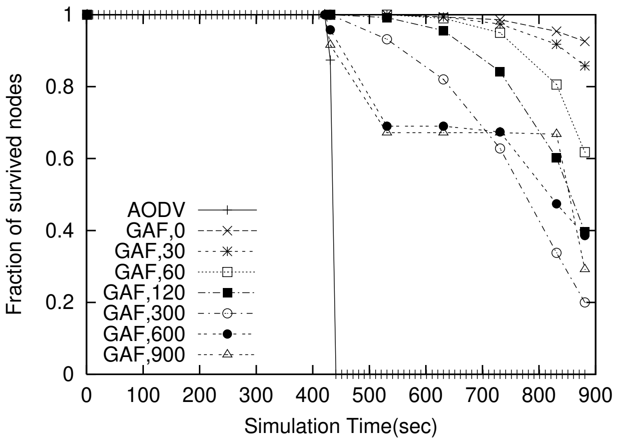

(a) GAF-b

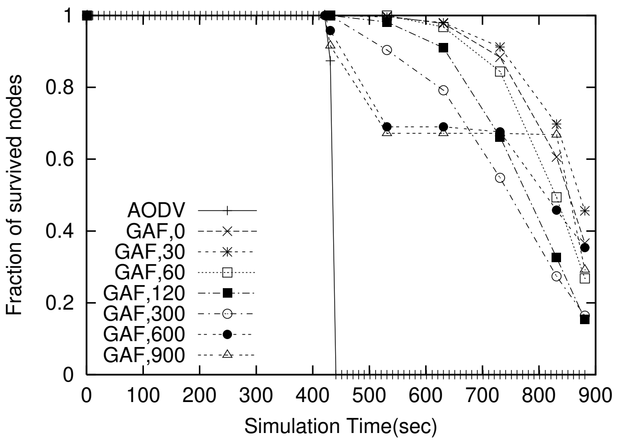

(b) GAF-ma

Figure 7: Network lifetime comparison: GAF-b,GAF-ma vs. AODV at high node speed (20m/s) under various pause time. Different traffic load does not affect the result. In the legend, "GAF,x" means running GAF/AODV with pause time x.

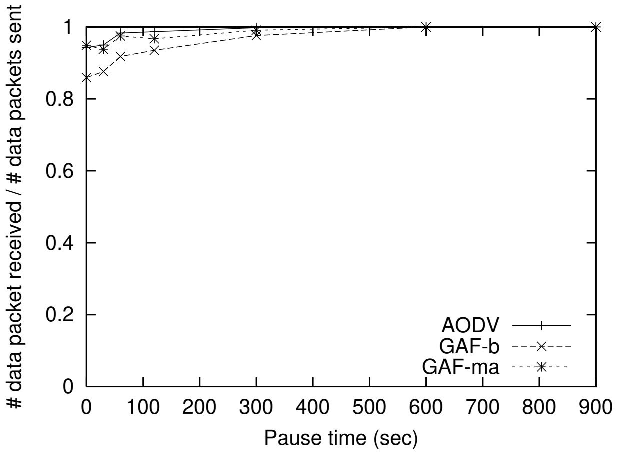

(a) Data delivery ratio

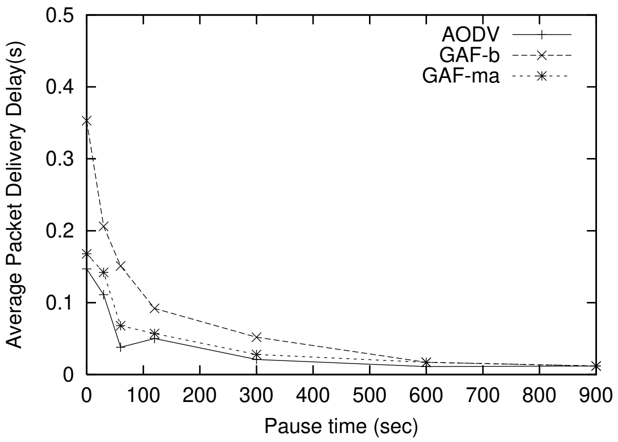

(b) Average delay

Figure 8: Data delivery as a function of pause time comparison: GAF-b,ma vs. AODV under moving speed 20m/s. Traffic load is 10pkts/s. Other load does not change the result.

network lifetime. In this section, we use simulations to quantitatively find how network lifetime changes with different node densities.

We address this issue by introducing more nodes into the same simulation scenarios used above. We changed the node number from 50 to 100 and 200 while keeping the area constant. Given our nominal radio range, this corresponds to 21, 42, or 84 nodes in range (See *ninra* definition later at Equation 6). There are 45 virtual grids in this area, so the mean number of nodes per grid will go from about 1.1 to 2.2 and 4.4. In order to have enough time to let all nodes to expire we extend the simulation time to 3600s. We consider only pause times of 0s and 3600s since in Section 3.3 we showed that pause times of 0s and 900s provide lower and upper bounds for network lifetime, so we consider only those here. We have evaluated both AODV and DSR at 50 nodes scenario and find similar behavior. Therefore we only report AODV here.

Our goal is to extend the lifetime of the network as a whole through energy conservation. Ultimately, the application wants to know how long the network can deliver information for it. In our simulation, our traffic keeps constant. We therefore can measure how long the network can deliver information by periodically monitoring the packet delivery ratio for every fixed time frame. We can also compare the monitored packet deliver ratio with AODV packet delivery ratio in its normal lifetime. If the monitored packet delivery ratio drops dramatically after time t, we can say that the network lifetime ends at time t. In order to quantitatively find the t, we can arbitrarily define that if the monitored packet delivery ratio is 20% less than the packet delivery ratio of AODV, we make a cut there. In our experiment, we monitor the packet delivery ratio every 100s in the extended network lifetime.

It is also interesting to know how node survives in the extended network lifetime. We demonstrate the node survivability by measuring node survival rates as a function of time. Note that the decrease of the number of nodes will affect routing fidelity. Normally the less survived nodes, the worse packet delivery ratio. However, we can't simply decide a network lifetime by the fraction of survived nodes. The reason is that in high node density, a small fraction of survived nodes may still be a big number of nodes which will deliver the traffic accordingly without losing routing fidelity.

Figure 9(a) shows our result of monitoring packet delivery ratio for every 100s for GAF-b and GAF-ma in a scenario of 100 nodes simulation with pause time 0, node speed 20m/s and traffic load 10pkts/s. We also mark the point where the packet delivery ratio drops more than 20% of AODV packet delivery ratio. For comparison, we put AODV packet delivery ratio inside the figure too.

From Figure 9(a), we find that roughly GAF-b quadruples AODV lifetime and GAF-ma triples the AODV lifetime. Before the drops, the packet delivery ratio remains almost constantly compared to AODV packet delivery ratio. This shows that both GAF-b and GAF-ma are very stable in delivering data during their lifetime. Such stability also shows that the number of active nodes is maintained at a constant

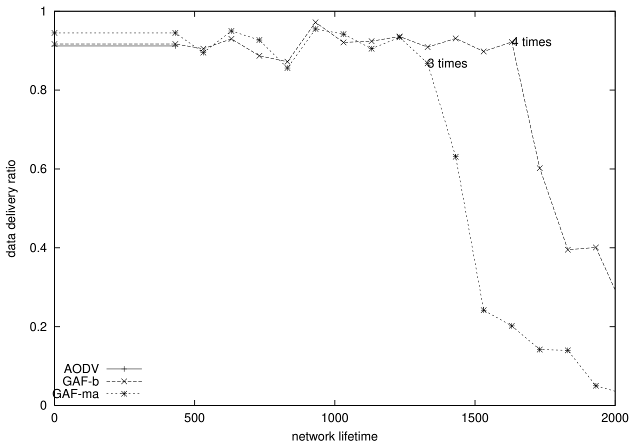

(a) packet ratio over network lifetime

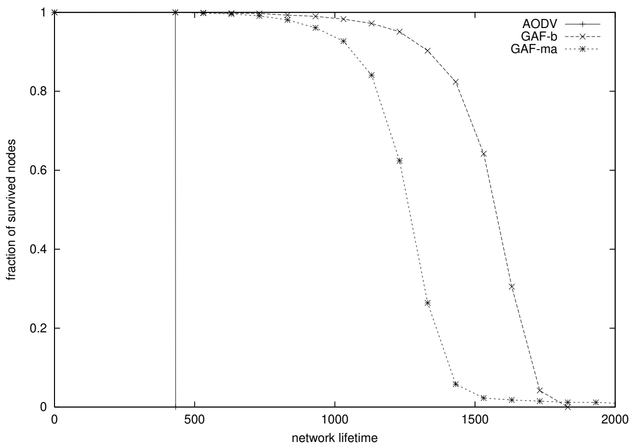

(b) fraction of survived nodes over network lifetime

Figure 9: Quantifying network lifetime over AODV. Simulating 100 nodes, at pause time 0, speed 20m/s, traffic load 10pkts/s. In figure, n times means n times of AODV lifetime.

level by GAF to maintain routing fidelity.

Node fraction over time is shown in Figure 9(b). We can find that the total number of survived nodes are maintained in the constant level during most of GAF-b and GAF-ma lifetime. This is mainly due to the load balance algorithm used in GAF so that all nodes try to evenly share their lifetime.

It is interesting to notice that from Figure 9(a), GAF-b has a very good data delivery quality in high node density. This is because the high node density provides more nodes to cover should an active node leaves the grid. Thus at high node densities, GAF-b may be suitable even at higher movement rates.

We next performed similar analysis for network lifetime with 200 nodes. These results are similar in shape to Figure 9

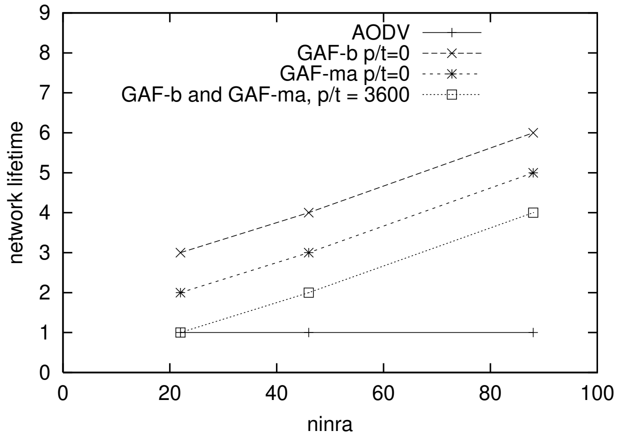

Figure 10: Network lifetime changes along with node density for GAF-b with constant mobility (pause time 0), GAF-ma with constant mobility, GAF with no mobility, and AODV.

and so are not presented here. Instead, we summarize the scalability of GAF as density increases in Figure 10. The y-axis shows network lifetime as defined above in multiples of AODV lifetime. To show the results independent of the number of nodes and the size of topography and radio, we define the *nodes in nominal radio area*, or *ninra*, as

$$ninra = \frac{\pi * R^2 * n}{w * h} \tag{6}$$

where R is the nominal radio range, n is the total number of nodes, and w and h are the width and length of the simulation area.

Figure 10 shows that GAF extends the network lifetime proportionally to the increase of node density regardless the mobility pattern while AODV's lifetime keeps flat for all scenarios. The greater node density can be used to increase network lifetime for about 4 to 6 times (depending on the mobility pattern) with a four-fold increase in density. Under high mobility, the saving reaches the upper bound of 6 times of network lifetime while it only reaches its lower bound, 4 times of AODV network lifetime, when all nodes keep still. (As described in Section 3.3, mobility allows longer lifetimes.)

3.8 Result sensitivity to propagation model

So far we have shown that GAF can provide longer lifetime with minimal loss in data delivery rates, and that these benefits are proportional to node density. However, these simulation studies considered a deterministic propagation model (the tworayground model). In reality, radio propagation is strongly affected by multi-path effects (fading). (These models are described in detail in wireless textbooks such as Rappaport [26].) We next revisit our simulations using a shadowing model to see if the results hold.

The shadowing model consists of two parts. The first one is

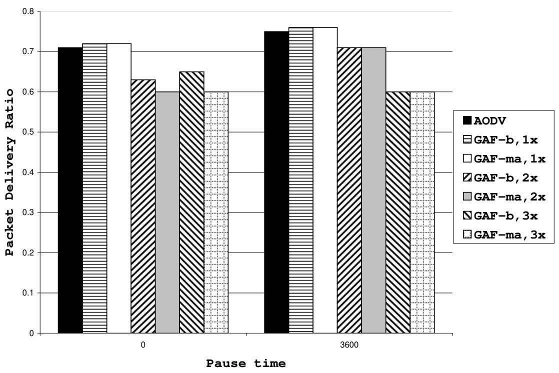

Figure 11: Comparison of packet delivery ratio for AODV, GAF-b and GAF-ma under shadowing model. GAF-*,nX means GAF-b or -ma at n times of AODV lifetime.

known as path loss model, which predicts the mean received power at certain distance. The path loss is usually measured in dB. The second part of the shadowing model reflects the variation of the received power at certain distance. The shadowing model extends the ideal circle model to a richer statistic model: nodes can only probabilistically communicate when near the edge of the communication range.

We select a set of typical outdoor value with path loss exponent 3.0 and the shadowing deviation 4.0 in our simulation to find out how GAF works under shadowing model. Due to the path loss, we expect to have worse data delivery ratio. However, normal ad hoc routing is also affected by the path loss. Our goal is to match GAF packet loss ratio with that of AODV under shadowing model when network lifetime can be extended.

We repeated our simulation by just replacing two-ray ground propagation model with shadowing propagation model. We summarized the result in Figure 11.

The figure shows that for AODV, compared with the simulation with tworayground propagation model, the packet delivery ratio is about 20% worse under shadowing model. GAF shows the the same downgrade. However, under shadowing model, GAF still can match AODV packet delivery ratio within AODV lifetime. With extended lifetime, GAF-b can maintain less than 10% packet delivery ratio decrease within 3 times of AODV network lifetime while GAF-ma can match the result within 2 times of AODV network lifetime. The result is the same as Figure 10.

The result shows that shadowing model does not change the result on GAF which we showed in the section 3.7. The shadowing model does affect the packet delivery ratio but its impact is the same on AODV and GAF.

3.9 Result sensitivity to location error

Many kinds of localization systems (including GPS) include inherent error. We would like to know how this error would affect GAF performance. We do not believe GAF will be affected by moderate location error or by large correlated error because of how GAF uses location information.

First, GAF will be robust to correlated error because it cares only about relative node position when assigning nodes to grids. If all nodes are accidently shifted in the *same* direction grid assignment will be consistently affected. The intentional error introduced in GPS is correlated (for nodes listening to the same satellites), so GAF should be robust to GPS error.

We have already shown that GAF is reasonably robust to somewhat inaccurate position information the scenarios with high mobility (20m/s speeds or 0 pause times). In addition, the GAF algorithm does not depend on accurate location information. For example, when a node A receives a discovery message from a node B, node B may have left the grid where node A resides but the discovery message sent by node B may still shows that node B is within the grid where node A belongs to. Our simulation has shown such inaccuracy does not affect GAF performance.

Finally, we artificially introduced location error in simulations. We modeled error by randomly recoding the location of each node in the range $[x - e, x + e]$ and $[y - e, y + e]$ for e of 5 meters. We then repeat our first phase simulation with this error; nodes do use these incorrect positions in place of their real locations for the GAF algorithm. We study the same methodology used in Sections 3.3, 3.4, and 3.5 to calculate the network lifetime, energy saving and data delivery quality. We found no substantial variation in these metrics. For example, in terms of the data delivery ratio under low mobility, they both reach 99%. Under high mobility, GAF-b has 85% data delivery ratio and GAF-ma has 95% data delivery ratio. The same is identical to that in section 3.6.

With uncorrelated error approaching the virtual grid size GAF will fail to work. We believe that we could accommodate such error by artificially reducing grid size, but exploring of this is future work.

4. RELATED WORK

Our work builds on related work for radio energy models, ad hoc routing protocols, and energy-aware MAC and application-level protocols.

Ad hoc routing protocols: A number of routing protocols have been proposed to provide multi-hop communication in wireless, ad hoc networks [22, 4, 23, 21]. Traditionally these protocols are evaluated in terms of packet loss rates, routing message overhead, and route length [5, 15, 10]. We evaluate our protocols by these metrics for comparison, and we add measures of power consumption and network lifetime to consider power consumption as well.

Both Chang and Tassiulas [7] and Pottie et al. [24] have recently suggested that one might select routes in an ad hoc network based on available energy. The effect of this work would be longer network lifetime. Our approach is to conserve energy by powering radios off rather than managing a fixed energy consumption, so our work complements their effort.

Heinzelman et al. present a set of protocols for communication in sensor networks based on flooding [14]. They examine the energy consumption of these protocols and show that suppressing duplicate transmissions of the same data can save power as calculated from a simple energy model (not considering energy consumption while radios are idle). Unlike their work, we consider more accurate power models and ad hoc routing protocols rather than flooding. These differences result in much different optimizations. We also consider optimizations based on adaptive fidelity that are specific to dense networks.

Cluster-based ad-hoc routing protocols: Some ad-hoc network routing approaches [8, 18] employ a cluster-based philosophy. They structure the ad-hoc network as a two-level network: in the lower level, nodes in geographical proximity create peer-to-peer networks. In each one of these lower-level networks, at least one node is designated to serve as a "gateway" to the higher tier. These gateway nodes create the higher-level network. Routing between nodes that belong to different lower-level networks is through the gateway nodes. Clustering schemes are often used with a TDMA MAC protocol to reduce the cost of radio listening.

Cluster-based approaches do not directly address the energy dissipation issues due to the cost of radio listening and overhearing, although a TDMA protocol helps some. In addition, they introduce the new cost of hierarchy formation, and possibly additional cost by forcing communication through gateways even for nodes that could directly hear each other.

Although GAF finds redundant nodes within each virtual grid, GAF is not a cluster-based algorithm. By using application-level information GAF can have much lower duty cycles even than clustering with a TDMA MAC. GAF also does not force the packets from the source nodes to go through all representative nodes in each grids in order to reach the destination although it is up to the ad hoc routing protocols to find the path.

Geographic Ad hoc routing: Ko and Vaidya presented a Location-aided Routing (LAR) approach to utilize location information to improve performance of routing protocols in ad-hoc networks [17]. They use location information to decrease overhead of routing discovery by limiting the search space for a desired route.

GAF is compatible with LAR as it is with other ad hoc routing protocols. GAF's approaches to energy savings are also complementary with those in LAR; LAR optimizations does not aim to reduce energy dissipation.

Grid location service (GLS) [20] provides location services by duplicating a node's location in a small subset of other nodes. With the help of GLS, a node can geographically forward data to the destination with the destination location. A node chooses its location servers in a hierarchy of grids with increasing size. Such grid-based partition is very close to the grid used in GAF. Since GAF is not a geographical routing protocol, there is no need for GAF to propagate location information to other nodes. Grid system with geographic information has different use in GLS and GAF: it helps choosing location servers in GLS while it helps finding node coverage in GAF.

Energy-aware MAC protocols: PAMAS is a MAC-level protocol where radios power off when not actively transmitting or receiving packets [28, 29]. PAMAS avoids the overhearing problem we discuss in Section 1, but it does not address the problem of energy consumption when nodes are idle. Solutions to overhearing are relevant, but for radios with high idle power consumption work such as we propose will be necessary.

TDMA protocols have been proposed to reduce energy consumption in sensor networks [24]. By reducing the duty cycle these protocols can trade idle-time energy consumption for latency. We believe TDMA MAC protocols will very important for power-constrained networks. Although we have not yet examined use of our approaches over TDMA protocols, our use of application-level information and node density can further improve power conservation.

Ram Ramanathan and Regina Rosales-Hain investigated the problem of creating a desired problem by adjusting transmit power in ad hoc network [25]. Ideally, the topology control can reduce energy use in dense network by controlling the number of neighbors of each node, and improve connectivity in a sparse network. Its two mobile protocols, LINT and LILT, use similar approach in AFECA [32] to collect neighbor information collecting by routing protocols and attempt to keep some level of number of neighbors of each node. While we believe that the transmit power control can adaptively conserve energy (presumably the idle energy cost can be reduced too) if the hardware provides enough support, we are concerned with the impact on the ad hoc routing protocols. Intuitively, the hop-by-hop transmission with restrained transmit power may lead to fragile route due to mobility. The paper does not give packet drops comparison with existing ad hoc routings, nor does the comparison under different mobility model.

IEEE 802.11 [19] supports ad hoc network configuration: mobile nodes are brought together to form a network on the fly. IEEE 802.11 also provides power management controls to allow disabling the transceiver to conserve energy. Although they specify how to turn off the radio, they do not discuss specific policies. We propose these policies assuming the presence of 802.11-like controls for basic and adaptive cases.

PicoNet: PicoNet proposes an integrated design of radios, small, battery powered nodes, and MAC and application protocols that minimize power consumption [2]. They reduce power consumption with a very low, application-dependent duty cycle (their paper does not specify, but presentations suggest intermittent polling with periods of 50 to 100s of seconds). They primarily use local base stations instead of multi-hop wireless routing, and assume frequent or continuous node movement. Their approaches are promising, but we are not aware of a detailed study if PicoNet power consumption. Our work differs from theirs by building on existing ad hoc routing protocols and by making use of adaptive fidelity to reduce power in dense node configurations.

Other examples of adaptive fidelity: We have previously described two algorithms that apply adaptive fidelity to ad hoc routing [32]. Unlike GAF, both these protocols require integration with the underlying ad hoc routing protocol. BECA powers of nodes independent of density, while AFECA samples traffic to evaluate node density and implement adaptive fidelity. We have found that with location information, GAF is able to provide much better energy savings that either BECA or AFECA.

SPAN [9] has the same goal as GAF and AFECA, to conserve energy and increase system lifetime by turning off redundant nodes without substantially affecting the connectivity of the network. Each SPAN node decides whether to sleep or join the backbone based on local topology information. SPAN selects coordinators using network topology information provided by a geographically-informed ad hoc routing protocol. In this sense, both SPAN and GAF take advantage of geographic information, however SPAN does so indirectly.

Bulusu et al. [6] are investigating the use of adaptive fidelity for localization systems. Given a field of energy-constrained beacon nodes, she is examining what duty cycle optimizes network lifetime.

5. FUTURE WORK

We have identified a number of areas for future work. Although we have studied GAF under several loads and scenarios, additional sensitivity analysis of the results is desirable. We would like to evaluate GAF sensitivity to mobility models other than random. We studied GAF under a relatively small number of traffic nodes; future work might consider heavier traffic. Larger ranges of location error are probable for some localization schemes so studies of larger error are appropriate. On the other hand, we are exploring algorithms to identify node redundancy without requiring location information. Finally, we only use a set of parameters for shadowing models, different shadowing models are also possible.

Additional exploration of network behavior as nodes fail is important. When does the network partition? How do mixes of nodes with different power affect the results? These results are likely sensitive to traffic mix.

Finally, experimentation is needed to validate these results with physical hardware in actual scenarios.

Following AFECA [32] this is the second adaptive fidelity algorithm we've applied to ad hoc routing. More work is required exploring the concept of adaptive fidelity in other contexts. For example, we would like to understand the performance difference of sensing in dense sensor networks when sensors are only enabled with some duty cycle.

6. CONCLUSIONS

We have demonstrated our approach to energy conservation for ad hoc routing. Power consumption in current wireless networks is idle-time dominated, so GAF focus on turning the radio off as much as possible.

GAF adapts sleep time based on node location scaling back node duty cycles (and so reducing routing "fidelity") when many interchangeable nodes are present. We have shown that it performs at least as well as an normal ad hoc routing protocol for packet loss and route latency, and yet it can substantially conserve energy, allowing network lifetime to increase in proportion to node density. This simple "add more to improve service" behavior of GAF and adaptive fidelity are particularly important as the numbers of embedded devices and the ratio of devices to humans increases.

7. ACKNOWLEDGMENTS

The authors would like to thank the CMU Monarch research group which provides the ad hoc network extension for ns-2 and aiding integration. We also thank Charles Perkins, Samir R. Das and Mahesh Marina for providing their revised AODV implementation to us. Dr. Samir R. Das and Mahesh Marina answers our questions about the details of operation of AODV.

8. REFERENCES

[1] 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.

[2] 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.

[3] 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`.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] Y.-L. Chang and C.-C. Hsu. Routing in wireless/mobile ad-hoc networks via dynamic group construction. *Mobile Networks and Application*, pages 27–37, May 2000.

[9] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris. Span : An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. In *To appear Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, page To appear, July, 2001.

[10] 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.

[11] 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.

[12] 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.

[13] S. Floyd, V. Jacobson, C. Liu, S. McCanne, and L. Zhang. A reliable multicast framework for light-weight sessions and application level framing. In *Proceedings of the ACM SIGCOMM*, pages 342–356, Cambridge, Massachusetts, Aug. 1995. ACM.

[14] W. R. Heinzelman, H. Kulik, and H. Balakrishnan. Adaptive protocols for information dissemination in wireless sensor networks. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 175–185, August, 1999.

[15] 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.

[16] 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.

[17] Y.-B. Ko and N. Vaidya. Location-aided routing (LAR) in mobile ad hoc networks. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 66–75. ACM, 1998.

[18] P. Krishna, N. Vaidya, M. Chatterjee, and D. Pradhan. A cluster-based approach for routing in dynamic networks. *ACM SIGCOMM Computer Communications Review*, pages 372–378, 1997.

[19] LAN MAN Standards Committee of the IEEE Computer Society. *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification*, 1999. IEEE Std 802.11, 1999 Edition.

[20] J. Li, J. Jannotti, D. D. Couto, D. Karger, and R. Morris. A scalable location service for geographic ad hoc routing. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 120–130, August, 2000.

[21] 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.

[22] 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 progrss.

[23] 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`.

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

[25] 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.

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

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

[28] 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.

[29] 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.

[30] 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.

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

[32] 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.