

Energy-Aware On-Demand Routing for Mobile Ad Hoc Networks

Nishant Gupta
OPNET Technologies, Inc.
7255 Woodmont Avenue
Bethesda, MD 20814
U.S.A.
and
Samir R. Das
Computer Science Department
SUNY at Stony Brook
Stony Brook, NY 11794-4400
U.S.A.
Contact Email: samir@cs.sunysb.edu

Abstract

On demand routing protocols for ad hoc networks discover and maintain routes on a reactive, “as-needed” basis. These protocols are attractive for their low routing overheads. We develop a technique to make these protocols energy-aware in order to increase the operational lifetime of an ad hoc network where nodes are operating on battery power alone and batteries cannot be recharged. Our techniques uses a new routing cost metric which is a function of the remaining battery level in each node on a route and the number of neighbors of this node. The idea of the cost metric is to be able to route around the nodes that are running low in battery for which alternate routes are available. In addition, rerouting is done proactively when any node *en route* starts running low on battery while the route is being actively used. Further, we save energy by switching off the radio interfaces dynamically during the periods when the nodes are idle. Simulation results using AODV protocol show that combination of these techniques results in a significant improvement of the energy budget of the network as a whole resulting in increased operational life time. Generally speaking, the improvement is higher for higher traffic loads and modest mobilities.

1 Introduction

A mobile ad hoc network (or MANET) [10] is a group of mobile, wireless nodes which cooperatively and spontaneously form a network independent of any fixed infrastructure (e.g., base stations or access points) or centralized administration. A node communicates directly with the nodes within radio range and indirectly with all others using a dynamically-determined multi-hop route. Though the major motivation of studying ad hoc networks comes from military usage, they will also be useful in any form of tactical communications such as disaster recovery, explorations, law enforcements, and in various forms of home and personal area networks, as well as sensor networks.

The MANET environment is typically characterized by energy-constrained nodes, variable-capacity, bandwidth-constrained wireless links and dynamic topology, leading to frequent and unpredictable connectivity changes. Many dynamic routing protocols MANET have been proposed and evaluated, some within the IETF MANET Working Group [10]. Of particular interest is the new class of on-demand, source-initiated protocols that set up and maintain routes on an “as needed” basis to reduce routing overheads. This approach is in sharp contrast with the traditional, proactive and shortest-path based protocols (e.g., link state and distance vector) that have been successfully used in wired networks, including the Internet.

The majority of the work reported in the literature focuses on the protocol design and performance evaluation in terms of traditional metrics such throughput, delay and routing overhead. However, much less attention has been paid in making the routing protocol energy efficient. In critical environments such as military or rescue operations, where ad hoc networks will be typically used, conserving of battery power will be vital in order to make the network operational for long durations. Recharging or replacing batteries will often not be possible. This makes the study in energy-aware routing critical. The challenge in ad hoc networks is that even if a host does not communicate on its own, it still frequently forwards data and routing packets for others, which drains its battery. Switching off a non-communicating node to conserve battery power may not be always a good idea, as it may partition the network.

Conventional on-demand routing protocols such as AODV [11, 12] and DSR [8] are energy-unaware. Routing is done based on shortest path, the cost metric either considers number of hops or end-to-end delay at the time when route is established. The protocols do not proactively modify routes until they break. If nodes are energy-constrained, such metrics may have adverse effect on the network lifetime on the whole. For example, a node that lies on several routes will die prematurely and the network may get partitioned. Since recharging or replacing the battery is not feasible in most of the ad hoc network applications, it is imperative to study and design routing protocols which are able to conserve node energy to prevent such premature death.

Our work focuses on augmenting the existing on-demand routing protocols and making them energy conserving. On-demand protocols are more suitable for this study as they typically have lower routing overhead than proactive, distributed shortest path protocols and thus have a low baseline energy consumption. We have used AODV as the base on-demand routing protocol. The techniques implemented are generic in nature and should be applicable to other on-demand routing protocols, such as DSR.

The rest of the paper is organized as follows. In the next section, we briefly describe the AODV protocol. In section 3, we describe our energy aware techniques as an extension of AODV. In section 4 we evaluate the performance of our energy-aware techniques vis--vis the baseline AODV via simulations. In section 5 we review related work. We conclude in section 6.

2 AODV Protocol Description

Ad hoc On Demand Distance Vector (AODV)[11, 12] is source initiated, reactive protocol. It discovers and maintains routes only if and when necessary. Route discovery works as follows. When the source requires a path to a particular destination, it broadcasts a route request (RREQ) packet in the ad hoc network. Nodes receiving RREQ record a *reverse* route back towards the source, using the node from which the RREQ was received as the next-hop, and then re-broadcasts the RREQ. If the same RREQ

is received more than once (via different routes), it is ignored. This way the RREQ packets is flooded to every node in the connected part of the network.




Figure 1: Reverse and forward route formation in AODV

When the RREQ packet reaches the destination, it sends a route reply (RREP) packet back to the source, using the reverse route. If an intermediate node has an up-to-date route to the destination, it may also send a RREP packet back to the source on behalf of the destination. As the RREP packet follows the path back to the source, the corresponding *forward* route is created at each intermediate node towards the destination (see Figure 2). Once the RREP packet reaches the source, data traffic can now flow along this forward route.

To prevent routing loops, AODV maintains a sequence number on each node. Any routing information transmitted on routing packets or maintained on a node is tagged with the last known sequence number for the destination of the route. AODV protocol guarantees the invariant that the destination sequence numbers in the routing table entries on the nodes along a valid route are always monotonically increasing. Other than preventing loops, sequence numbers also ensure freshness of routes. Given a choice of multiple routes, the one with a newer sequence number is always chosen.

An important feature of AODV is maintenance of timer based states in each node, regarding utilization of individual routes. A route is “expired” if not used recently. A set of predecessor nodes is maintained for each routing table entry, indicating a set of neighboring nodes that use that entry to route data packets. These nodes are notified with route error (RERR) packets when the next hop link breaks. Each predecessor node, in turn, forwards the RERR to its own set of predecessors, thus effectively erasing all routes using the broken link. This RERR is thus propagated to each source routing traffic through the failed link, causing the route discovery process to be reinitiated if routes are still needed.

3 Energy Aware Routing in AODV

We take a two pronged approach for conserving the power budget of individual nodes. In the first, we modify the routing protocol to route around nodes with lower power budgets. In the second, we strategically turn off the radio interfaces to conserve energy further. The techniques are orthogonal to each other and are described separately in the following subsections.

3.1 Energy Related Cost Metric

On-demand protocols such as AODV typically pick the shortest path route during the route discovery process, and then sticks to this route until it breaks. Continuous use of the route may drain the nodes of battery power. This is particularly true if one or more nodes are on other routes as well. Note that each message transmission and reception drain battery power. If a node runs out of battery energy and unable to forward any messages, it effectively falls out of the network. In this case, the route breaks and AODV finds an alternate route via another route discovery. However, nodes dying such as this adversely affect the operational life time of ad hoc network. First, many applications where the dying nodes are the communication end points will fail. Second, even when the dying nodes are not the communication end points, network connectivity will become sparser and network partition becomes more likely.

The goal of our protocol is routing or re-routing around nodes low on battery power as far as possible. This will prolong the network lifetime. However, this should be done in such a way that other useful performance metrics (e.g., end-to-end delay and throughput) are not compromised in a significant way. We take a two-step approach to design the adaptive energy-aware protocol. First, the nodes are classified according to their remaining battery energy. Depending on their classification the nodes react differently to the routing protocol dynamics. Second, a new cost function is used as routing metric taking into consideration both the hop-wise distance and the battery levels of the nodes.

The nodes are classified according to the following energy zones.

- **Normal Zone:** Node are in the normal zone if their current energy level is greater than 20% of their initial energy. This signifies that the number of hops should be the deciding factor in determining the cost of routing data packets through these nodes as have ample energy at their disposal.
- **Warning Zone:** Nodes are in the warning zone if their current energy level lies between 10–20% of the initial energy. This signifies that the nodes are running low on energy and the protocol should avoid the use of these nodes if possible.
- **Danger Zone:** Here, the nodes have less than 10% of their initial energy. This means that the nodes are really low on the battery level and should only be used if there is no other cost-effective alternative.

The percentages of the initial energy reserves used in the definitions are somewhat arbitrary, and do not need to be exact. Figures in the same ballpark should work as well. It is assumed that the initial energy of the node is the maximum energy provided by the battery when it is fully charged. The purpose of assigning zones to nodes with various battery levels is to assign different costs for routing via nodes in different zones. The cost of routing a data packet through nodes in Warning (Danger) Zone is higher than the cost involved with the nodes in the Normal (Warning) Zone. This is to encourage the route discovery mechanism to explore alternate routes with higher battery power.

The total cost of a n -hop route using nodes x_1, x_2, \dots, x_n is given by $Cost = \sum_{i=1}^n c(x_i)$, where, the cost function for each node $c(x_i) = C_{normal}, C_{warning}$ or C_{danger} , depending on the zone of the node. $C_{normal}, C_{warning}$ or C_{danger} are predefined costs such that $C_{normal} < C_{warning} < C_{danger}$.

As an improvement to the suggested scheme, we keep track of the number of neighbors for each node from the neighborhood information gathered from the routing protocol. If the node lies in the Warning or Danger Zone and has sufficiently large number of neighbors then, we *increase* the cost of

routing of data packets through that node by a factor proportional to the number of neighbors. This helps in further avoiding a node with low battery levels if it has nodes in its vicinity that can do the same work. Now the new cost function for each node is $c(x_i) = C_{normal}, k_i C_{warning}$ or $k_i C_{danger}$, where k_i is proportional to the number of neighbors of node x_i .

3.2 Protocol Modifications Using Energy Cost

The above cost metric is used in AODV route discovery. Each RREQ packet flooded in the network builds up the cost for the path traversed so far by the packet. Each routing table entry also maintains the cost for that route. In regular AODV any node acts on only the very first RREQ received per route discovery flood. Duplicates of the RREQ received via alternate routes are ignored. However, use of this new cost metric requires that AODV acts on all such duplicates if they carry a lower cost metric. If a RREQ arrives with a lower cost metric (compared the metric in the routing table entry for the source indicating the cost of the reverse path), it is forwarded if the node is not the destination and does not have a route to the destination; otherwise it is replied to.

In AODV routing activity is reactive. It is possible that once a route is set, it remains active for a long period of time. In such cases, it might happen that one or more nodes on the route may move from one energy zone to another as they deplete their battery power in forwarding data packets. If this continues for a long time then some nodes may die. To ensure that the route is recalculated when the battery level depletes sufficiently to move any node on an active route into a different zone, such a node sends a route warning (RWARN) packet back to the source(s) using that route. The warning packet is propagated much like RERR, except that the route is *not* erased. Thus the flow of the data packets is not interrupted. A new route discovery process is initiated at the source on receipt of RWARN.

The new route discovery process do *not* selectively ignore the node(s) that sent RWARN. However, now such nodes incur higher cost using the method above. If a less expensive route is found, the routing tables of the appropriate nodes automatically switch to the new route. If no less expensive route is found, the old route continues to be used.

3.3 Optimizations at the Network Interface

Some analysis with our simulation models that will be presented momentarily has shown that nodes in the ad hoc network spends a large fraction of time being idle, that is, not transmitting or receiving packets. The main reason is that the nodes not on any active route rarely perform any communication, except when they need to propagate route request floods. Note that inactive routes are expired in AODV based on a timer. Several studies that performed measurements on network interface indicated that switching off the network interface incurs substantial energy savings.

In order to develop a technique to derive energy savings from the network interface we assume that the node is either in *active* or *sleep* state. Generally speaking, a node is in *active* state as long as it has an active route and it has packets in its buffer which needs to be sent, and otherwise it is in *sleep* state. If the node is receiving/transmitting a packet, it waits for a period of T_{active} after it has completed the communication activities. If no more packets are transmitted or received the node switches to the *sleep* state. In sleep state, the radio interface is switched off and the node is unable to transmit or receive any packet. The node switches back to *active* state after a period of T_{sleep} in the sleep state, and remain there for at least T_{active} time. If no packet is transmitted or received it switches back to sleep again. While sleeping, the node can become active earlier than T_{sleep} , if it generates any packet destined for others. Needless to say, the radio interface is switched on when the node is in *active state*.

Note also that regardless of the state, the node itself is always powered on; only the radio interface is switched on and off. Though powering off the node as a whole in the *idle* period is possible, in this paper we only focus on the networking components, and do not investigate this possibility. The power budget of the node used in the simulations later is used solely by the network interfaces.

The fraction of time the interface is switched off depends on the value of T_{sleep} . Choice of this value presents a critical tradeoff. While larger values are always better for power savings, large sleep periods mean that a node may possibly miss route discovery floods while it is in sleep state. Thus routes cannot be formed via the sleeping nodes. This may result in formation of suboptimal routes. In some extreme cases, the only routes that can be formed will be via the sleeping nodes, and since they cannot participate, route discovery will fail. This will impact both delay and throughput. Thus T_{sleep} cannot be very large. T_{sleep} cannot be very small either as from practical perspectives some minimum settling time is needed for the radio interface switching between on and off states. According to the measurements made in [16] using the first generation, 915 MHz WaveLAN radios, around 100 ms must expire from the time the radio is turned on and to the time it is ready to send the first packet. Though much depend on the actual electronics of the radios, we use this measurement as guidance and set the sleep period T_{sleep} at 115 ms. Even with small T_{sleep} , large power savings are possible if both T_{active} and T_{sleep} are small.

3.4 Performance Tradeoffs

The increase in the lifetime of the network and savings in the energy does not come free, and presents interesting tradeoffs with other useful performance metrics. For example, use of energy cost metric for routing may sometimes increase the number of hops, thus increasing the end-to-end delay of data packets. Also, switching off radio interface can interfere with routing, as routes cannot be formed via sleeping nodes. This also can also result in longer routes or, in extreme cases, failure of route discovery. The latter may result in longer delays or lost packets due to buffer overflows at the source. The exact nature of this tradeoff can only be understood by detailed performance evaluation, which we present in the next section.

4 Performance Evaluation

We simulated our energy-aware routing techniques as an extension to AODV for a mobile ad hoc network. We use the *ns-2* simulator [5] to implement the protocols. The latest version of AODV protocol [12] is used. The energy model used bears similarities to earlier studies [6, 16]. It is assumed that that the radio interface, when powered on, consumes 1.15W when listening to the channel for any incoming packet, 1.2W while actually receiving a packet and 1.6W while transmitting a packet. These values correspond to direct measurements on 915 MHz WaveLAN radios as reported in [17]. The physical radio model used in the simulator is very similar to WaveLAN; so these are appropriate numbers to use. The cost functions for nodes in different battery zones chosen as: $C_{normal} = 1$, $C_{warning} = 3$, and $C_{danger} = 4$. Thus, cost of routing via a node in warning (danger) zone is three (four) times the cost of routing through a node in normal zone. The multiplicative factor k_i is chosen as the number of neighbors of node x_i divided by 5. These parameters are some what ad hoc at this point and these values have been found to work well in the scenarios tested. In our experience, several factors influence their choice including network size and average node density. The values of different timers are chosen as follows: $T_{active} = 100$ ms and $T_{sleep} = 115$ ms.

4.1 Experimental Scenarios and Performance Metrics

In our experiments 50 nodes move around in a rectangular area of 1500m X 300m according to a mobility model (random waypoint, as described in [1]). Each node uses IEEE 802.11 standard [4] MAC layer. The radio model is very similar to the first generation WaveLAN radios with nominal radio range of 250m. The nominal bit rate is 2 Mbps. In this mobility model each node moves towards a random destination and pauses for certain time after reaching the destination before moving again. In our simulations, the nodes move at an average speed of 20m/sec. The pause times are varied to simulate different degrees of mobility. The traffic sources start at random times towards the beginning of the simulation and stay active throughout. The sources are CBR (constant bit rate) and generates UDP packets at 4 packets/sec, each packet being 512 bytes. Note that very similar mobility and traffic models have been used in earlier simulation studies of ad hoc networks [1, 7, 3]. Each simulation is run for 900 seconds simulated time. Each point in the plotted results represents an average of three simulation runs with different random mobility scenarios.

The following performance metrics are evaluated. The first three metrics are typical metrics usually evaluated for analyze performance of routing protocols with best-effort traffic. The remaining metric are useful for evaluating the efficacy of energy-aware routing.

- *Packet delivery fraction*: measured as the ratio of the number of data packets delivered to the destination and the number of data packets sent by the source.
- *End-to-end delay*: measured as the average end-to-end latency of data packets.
- *Normalized routing load*: measured as the number of routing packets transmitted for each data packet delivered at the destination.
- *Battery energy*: measured as the total amount (in Joules) of remaining battery energy at the end of the experiment.

We also evaluate two more power-related metrics to measure how soon nodes are dying out of power and how many nodes are dead (i.e., have zero energy) at the end of the simulation. These metrics are useful in addition to total battery energy as they indirectly determine connectivity of the network and thus the useful lifetime.

4.2 Simulation Results

We have plotted the performance of the base AODV and energy-aware AODV for various metrics. Number of traffic sources and pause times are varied to reflect various loads and mobility. Note that pause time = 0 means constant movement and pause time = 900 sec means stationary network.

The initial energy for each node in this set of simulations is 1060 Joules, which represent a combined network wide initial energy of 53,000 Joules. Note that the initial energy was set at this value with some trial so that we can effectively demonstrate the difference in behavior of the base and energy-aware protocols. If the initial energy is too high, the energy-aware techniques based on battery levels will not kick in and essentially only the baseline routing protocol will be operational throughout the experiment. On the other hand, initializing with too low energy will make the nodes die too soon, and beyond this point the traditional metrics like delivery fraction and delay will be meaningless. We have found that in order to demonstrate effectively the behavior of different metrics, we need to set

the initial energy and simulation run-length combination in such a way that the nodes do have a small energy remaining (about 10–20% of the initial) at the end. Initializing with higher energy does not let the rerouting techniques kick in as the nodes remain the normal zone throughout. We, however, perform some experiments with lower initial energy to demonstrate how soon the nodes die in different scenarios and protocols.

In Figure 2 observe that the remaining energy at the end of the simulation is much higher for energy-aware AODV than the base AODV. For the chosen parameter values, the improvement is 7-8 times for low traffic (10 and 20 sources) and up to 30 times for high traffic (30 sources). Note, however, these factors can be a little deceiving as they depend strongly on the initial energy and the simulation run length. For example, the improvements may not be this substantial if run length is lower. On the other hand, a longer run length may produce a larger factors of improvement of the remaining energy. Very long run lengths will make all nodes die. This will be demonstrated later.


The next two figures indicate the regular performance metrics, *viz.* delivery fraction and delay. In Figure 3 notice that energy-aware AODV is at par with AODV for 20 and 30 sources, but delivers a marginally lower number of packets for 10 sources. Note that for low number of traffic sources, many nodes are sleeping sporadically. Such node may fail to respond routing activities. Thus some route discovery attempts may fail and some other may obtain longer routes avoiding the sleeping nodes. This increases both packet losses and packet delay. This problem reduces with higher number of sources as a relatively lower number of nodes sleeps with higher traffic diversity.

Similar arguments apply in the analysis of the delay plots in Figure 4. Notice that energy-aware routing presents a much higher delay for 10 sources, marginally higher delay for 20 sources and lower delay for 30 sources. The lower delay for 30 sources is attributed to frequent rerouting with leads to reduced congestion and better load balancing. With higher traffic load, the nodes depletes their energy faster, and thus rerouting is relatively more frequent. Notice also slightly improved delivery fraction for 30 sources in Figure 3 due to the same reason. The improvement is more pronounced for lower mobility as the base AODV hardly changes routes in such cases. Thus congestions do not go away automatically, but they do for energy-aware routing.


From Figure 5, we find that routing load increases for energy-aware AODV. The increase is justified as we are allowing more RREQs to be propagated for each route discovery and doing additional route discoveries via the new warning packets. The increase is substantial for smaller number of sources, but not significant for higher number of sources (roughly by a factor of 4–5 for 10 sources, a factor of 2 for 20 sources and a factor of 1.5 for 30 sources). The interesting thing to note is that overall energy savings are still possible even with a higher routing load and somewhat longer routes in energy-aware AODV. This savings are primarily due to the sleep mode introduced at the radio interfaces.

Figure 6 shows the number of nodes with zero remaining energy at the end of the simulation. Note that no nodes have zero energy for either protocols for 10 sources. This is because of low traffic load does not deplete the nodes' energy reserves sufficiently given the initial energy and simulation run length. However, node starts dying at large numbers for the base AODV for 20 and 30 sources, while energy-aware AODV demonstrates comparatively excellent performance with nodes never dying, except occasionally for the 30 node case.


To get an idea of the useful lifetime of the network, we initialized the nodes with lower energy (500J per node) and ran the simulation for the same 900 sec. Our interest here is to determine how soon the nodes start dying in various protocols. We keep track of the number of survived nodes (*i.e.*, nodes with non-zero energy) at every 10 sec interval. In Figure 7 (a) and (b) we show the results for low (900 sec pause time) and high (0 sec pause time) mobilities, respectively for 10, 20 and 30 sources.



(a) 10 sources

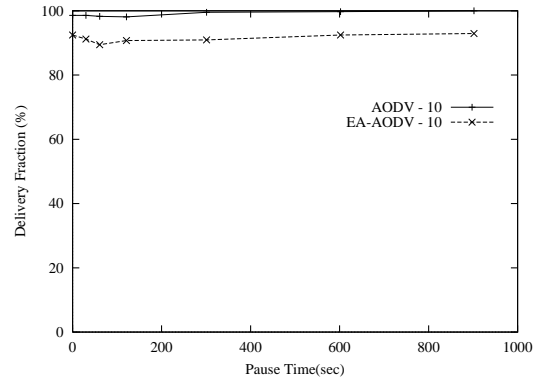


(b) 20 sources

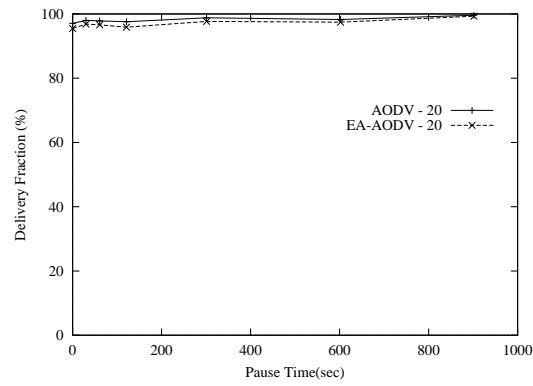


(c) 30 sources


Figure 2: Remaining energy for the 50 node model with various numbers of sources.



(a) 10 sources




(b) 20 sources




(c) 30 sources


Figure 3: Delivery fraction for the 50 node model with various numbers of sources.



(a) 10 sources




(b) 20 sources




(c) 30 sources


Figure 4: Average delay for the 50 node model with various numbers of sources.



(a) 10 sources




(b) 20 sources




(c) 30 sources


Figure 5: Normalized routing load for the 50 node model with various numbers of sources.



(a) 10 sources




(b) 20 sources




(c) 30 sources

Figure 6: Number of nodes with zero remaining energy for the 50 node model with various numbers of sources.



(a) 0 sec. pause time



(b) 900 sec. pause time

Figure 7: Number of survived nodes vs. simulation time for the 50 node model with various no. of sources.

Note that in base AODV all nodes die at about the same time. This is because even in the idle state the nodes consume a considerable amount of power. However, with our energy-aware techniques nodes die a little later and less drastically. This is particularly true for lower traffic situations (smaller no. of sources) and for lower mobility. We indeed performed experiments with intermediate mobilities and the results fall in between the two extremes presented here.

5 Related Work

The problem of trying to conserve battery usage within a mobile ad hoc network is not new. Previous other work has gone in different energy conserving strategies spanning different network layers. In [14] a power-aware multiple access protocol (PAMAS) was proposed. Here, a node turns off its radio interface for a specific duration of time, when it knows that it will not be able to send and receive packets during that time because of the possibility of multiple access interference. The sleep time is of the order of packet duration, which could be very small in modern wireless LANs like 802.11b (nominal bit rate 11 Mb/sec). This means that the radios must be capable of very fast transition between on and off periods (within a millisecond in 802.11b, e.g.). The state-of-the-art in commercial radios are far from this goal. For example, experimental studies in [16] indicate that it takes about 100ms for COTS radios like WaveLAN to come out of sleep state and be ready to transmit/receive the first packet. However, PAMAS approach would be quite viable for low-bandwidth wireless networks, in radios where small packets can be combined to form large packets, or in radios with fast settling periods. Our protocol do not have any dependency on the property of the electronics and would be able to gain further from PAMAS-like support in the MAC layer.

In [15], several energy-aware metrics were discussed that will result in energy-efficient routes. The metrics included maximizing the time to network partition and reducing variance in node power levels. It is hard to use these metrics directly in a network without any central control. However, our technique uses power balancing by rerouting around nodes with low energy reserves. This indirectly achieves the above goals. In [15] performance evaluations using small (20 nodes) random graph model showed improved per packet energy cost with a energy-aware cost model (cost of route is the sum of costs of each node on the route, where cost for a node is a linear or quadratic function of the energy expended thus far in the node) compared to shorest-path routing. Our cost model is similar though we use a piece-wise model dependent on the energy reserve of a node. Some of our observations are also in line with [15], such as the savings are greater in larger and denser networks, and also in moderate traffic situations as opposed to low or high traffic situations. In should be mentioned that [15] uses PAMAS in the MAC layer that contributed to greater savings.

In a more recent work [2], the network longevity was the overall goal in reducing the battery consumption. An algorithmic approach of a class of flow augmentation algorithms coupled with flow redirection was used. Unlike the conventional approach of minimizing the cost of the route from a source to a given destination, the strategy here was geared towards balancing the battery usage among the nodes in the network in proportion to their energy reserves. The algorithms were centralized; no deployable distributed solution was provided.

In [18], a power saving technique was proposed with AODV routing protocol. Here the nodes are made to sleep during the idle periods. Choice of sleep periods are somewhat different from ours. Mobile networks have not been considered and no routing cost model has been used. In general, powering radios off during periods a node is known not to take part in any critical communication is a popular technique for energy saving. PAMAS [14] does it in MAC layer; [18] does it in the routing

layer; transport layer solutions appear in [9]; some application-layer ideas can be found in [16]. In the latter two papers the solutions are not specific to ad hoc networks.

There has been a growing interest in transmit power control (see, for example, [13]) in the ad hoc networking community for reasons of power savings and to increase spatial reuse. Minimum transmit power necessary to successfully transmit a packet to a neighboring node follows a power law relation with the physical distance. Thus there can be interesting trade-offs between total energy needed to send a packet between a source-destination pair, delay (transmitting with lower power will increase number of hops and hence delay) and degree of network connectivity. Though significant potential exists with a power-controlled MAC, none of the COTS radios support power-control. We have not used transmit power control in our work.

The results obtained from the packet-level simulation are sensitive to the various parameters of the system. These parameters include the initial level of energy, the categorization of different energy zones, the cost associated with the various energy zones, the timings for the sleep and the idle periods and the extent to which active neighborhood effects the cost of routing the data packets. A detailed analysis of the effect of each of these parameters individually and the together with the rest of them on the results is necessary.

6 Conclusion

On-demand routing protocols are useful for mobile ad hoc network environment for their low routing overheads. However, if battery energy is not taken into consideration in their design, it may lead to premature depletion of some nodes' battery leading to early network partitioning. We have proposed adaptive energy-aware routing techniques as an extension to the AODV protocol that uses a new routing cost model to discourage the use of nodes running low on battery power. It also saves energy by turning off radios when the nodes are idle.

The results obtained from implementing these techniques are favorable and encouraging. Performance evaluation using a routing simulator shows that the longevity of the network can be extended by a significant amount. For the chosen simulation parameter set, in low traffic static scenarios this results in a increase in the lifetime of the network by up to 65%. For high traffic scenarios this increase ranges from 8-26%. The remaining energy is about 7-8 times more than the base AODV for low traffic situations and it goes up to 30 times for high traffic networks. There is a slight detrimental effect on delivery fraction and delay, except at high traffic scenarios, where both actually improves due to reduced congestion. This is an interesting positive side effect of the energy-aware protocol. Routing load, however, is consistently high, more at low traffic scenarios. Overall, we conclude that the energy-aware protocol demonstrates significant benefits at high traffic and not-so-high mobility scenarios. We expect that these scenarios will be common in ad hoc networking applications.

Even though we implemented the protocol on AODV, the technique used is very generic and can be used with any on-demand protocol. The energy-aware protocol works only in the routing layer and exploits only routing-specific information. It will be interesting to exploit MAC-layer specific information such as PAMAS, and transport and application-layer specific informations to optimize energy-usage further. We will study these avenues in our future research.

References

- [1] J. Broch, D. A. Maltz, D. B. Johnson, Y-C. Hu, and J. Jetcheva. A performance comparison of multi-hop wireless ad hoc network routing protocols. In *Proceedings of the 4th International Conference on Mobile Computing and Networking (ACM MOBICOM'98)*, pages 85–97, October 1998.
- [2] Jae-Hwan Chang and Leandros Tassiulas. Energy conserving routing in wireless ad-hoc routing. In *Proceedings of the IEEE INFOCOM 2000 Conference*, Tel Aviv, Israel, March 2000.
- [3] Samir R. Das, Charles E. Perkins, and Elizabeth M. Royer. Performance comparison of two on-demand routing protocols for ad hoc networks. In *Proceedings of the IEEE INFOCOM 2000 Conference*, pages 3–12, March 2000.
- [4] IEEE Standards Department. Wireless LAN medium access control (MAC) and physical layer (PHY) specifications, IEEE standard 802.11–1997, 1997.
- [5] Kevin Fall and Kannan Varadhan (Eds.). *ns* notes and documentation, 1999. available from <http://www-mash.cs.berkeley.edu/ns/>.
- [6] L.M. Feeney. An energy consumption model for performance analysis of routing protocols for mobile ad hoc networks, July 1999.
- [7] Per Johansson, Tony Larsson, Nicklas Hedman, and Bartosz Mielczarek. Routing protocols for mobile ad-hoc networks - a comparative performance analysis. In *Proceedings of the 5th International Conference on Mobile Computing and Networking (ACM MOBICOM'99)*, pages 195–206, August 1999.
- [8] D. Johnson and D. Maltz. Dynamic source routing in ad hoc wireless networks. In T. Imielinski and H. Korth, editors, *Mobile computing*, chapter 5. Kluwer Academic, 1996.
- [9] R. Kravets and P. Krishnan. Power management techniques for mobile communication. In *4th International Conference on Mobile Computing and Networking (ACM MOBICOMM'98)*, pages 157–168, October 1998.
- [10] J. Macker and S. Corson. Mobile ad hoc networks (MANET). <http://www.ietf.org/html.charters/manet-charter.html>, 1997. IETF Working Group Charter.
- [11] Charles Perkins and Elizabeth Royer. Ad hoc on-demand distance vector routing. In *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, pages 90–100, Feb 1999.
- [12] Charles Perkins, Elizabeth Royer, and Samir R. Das. Ad hoc on demand distance vector (AODV) routing. <http://www.ietf.org/internet-drafts/draft-ietf-manet-aodv-07.txt>, Nov 2000. IETF Internet Draft (work in progress).
- [13] Ram Ramanathan and Regina Rosales-Hain. Topology control in multihop wireless networks using transmit power adjustment. In *Proceedings of the IEEE INFOCOM 2000 Conference*, 2000.

- [14] S. Singh and C.S. Raghavendra. PAMAS - power aware multi-access protocol with signalling for ad hoc networks. *ACM Computer Communication Review (ACM CCR'98)*, July 1998.
- [15] S. Singh, M. Woo, and C.S. Raghavendra. Power-aware routing in mobile ad-hoc networks. In *ACM/IEEE International Conference on Mobile Computing and Networking*, pages 181–190, October 1998.
- [16] Mark Stemm and Randy Katz. Measuring and reducing energy consumption of network interfaces in hand-held devices. Technical report, UC Berkeley, August 1997.
- [17] Bruce Tuch. Development of WaveLAN, an ISM band wireless LAN. *AT&T Technical Journal*, 72(4):27–33, July/Aug 1993.
- [18] Ya Xu, John Heidemann, and Deborah Estrin. Adaptive energy-conserving routing for multihop ad hoc networks, October 2000.