

Elective Participation in Ad Hoc Networks Based on Energy Consumption

Marc R. Pearlman[‡], Jing Deng[†], Ben Liang[‡], and Zygmunt J. Haas[§]

School of Electrical and Computer Engineering, Cornell University, Ithaca, New York 14853, USA

Abstract—In ad hoc networks, each node utilizes its limited resources to carry out the collective operation of the network. It is not always in the best interests of the network’s nodes to demand the continuous participation of all nodes in the network operations. In this work, we propose an Energy Dependent Participation (EDP) scheme, where a node periodically re-evaluates its participation in the network based on the residual energy in its battery. More importantly, a node gives special consideration to supporting the communication needs of its active network applications and preventing further network partitioning. EDP’s localized partition checking algorithm is particularly well suited for the Zone Routing Protocol, where the link-state information is proactively maintained within each node’s local zone and routes to faraway nodes are reactively obtained via global queries. Through simulations, we evaluate the impact of our proposed scheme on battery life and network connectivity. Our results suggest that the EDP scheme can increase the usable lifetime of a battery-constrained ad hoc network by over 50%.

I. INTRODUCTION

Ad hoc networks are formed by a collection of mobile nodes that organize themselves into a functioning network without reliance on existing infrastructure. Each node is equipped with a low power radio transceiver, providing direct communication between a set of nearby neighbors. In order to provide network-wide communication, each node is called upon to serve as a router. The result is a distributed multi-hop network with a time varying topology.

Ad hoc network nodes are typically portable and thus powered by limited capacity batteries. Not surprisingly, energy conservation in this type of environment is an area of active research. In [1], [9], [10], and [14], the problem of determining the optimal transmission radius for minimal interference and energy consumption is considered. In [11], [9], [2], [8], and [5], the problem of minimum energy routing is addressed.

However, relatively little work has dealt with power consumption of nodes when they are active but not transmitting. Experimental work presented in [12] has suggested that the power consumption of receiving or idling nodes can be significant. Therefore, in order to more aggressively conserve energy, one must consider the possibility of reducing not only the transmission periods, but also the duration of receiving and idling states.

In this paper, we consider a strategy, called Energy Dependent Participation (EDP), that gradually reduces a node’s participation in network operations as the battery’s energy is consumed. We propose an efficient mechanism that prevents node withdrawal from the network, if that node’s absence would cause network partitioning. Furthermore, we show how the proposed scheme can work seamlessly with the Zone Routing Protocol [6] to significantly improve the *usable lifetime* of an energy constrained ad hoc network.

The rest of this paper is organized as follows. Section II summarizes related research work. Section III describes the radio transceiver power consumption model. Section IV presents the proposed energy dependent participation scheme. Section V describes the mechanisms for maintaining network applications and preventing network partitioning. Section VI discusses the relation between EDP and the routing protocol. Section VII presents the simulation results. Finally, Section VIII present the conclusions.

II. RELATED WORK

The approach of temporarily turning off devices has been considered in other research work. In [7], an indexing method is proposed for reducing the energy consumption of hand-held devices while accessing data from a broadcast channel. In this method, the server broadcasts its files interleaved with periodic indices to the files, and a hand-held device remains in doze mode most of time, only tuning in periodically in order to download the indices and the required data. This approach is suitable for the client-server environment. In contrast, the emphasis of our work is in conserving the energy of mobile terminals in the multi-hop and peer-to-peer ad hoc networks.

In [13], centralized algorithms were introduced to schedule the recharging of low energy nodes while maintaining network reliability. Centralized algorithms are generally not suitable for ad hoc networks due to the frequently changing packet-radio environment. The elective participation strategies that we propose are distributed and independently carried out by each individual node.

Two new schemes related to our work have been recently proposed in [3] and [15]. In [3], a set of “coordinators” are used to forward data packets. In this scheme, the selection of coordinators is mainly based on local connectivity, and the node residual energy comes into play only if there are multiple contending coordinator candidates. Furthermore, the scheme requires clock synchronization and the non-coordinators’ participation in maintaining local connectivity and clock synchronization by sending

[‡] Currently with the GE Global Research, Niskayuna, New York, USA.

[†] Currently with the New York State Center for Advanced Technology in Computer Applications and Software Engineering (CASE) and the Center for Systems Assurance (CSA), in the Department of Electrical Engineering and Computer Science (EECS) of Syracuse University, Syracuse, NY, USA.

[‡] Currently with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, Ontario, Canada.

[§] This work has been sponsored in part by the ONR contract no. N00014-00-1-0564, and the NSF grants no. ANI-9980521 and no. ANI-0081357.

and receiving beacons at a frequency pre-determined by the underlying MAC protocol. In our approach, the dominating factor for deciding whether or not a node participates in data forwarding is its residual energy. This energy-centric approach is easy to implement and does not require clock synchronization. In addition, the non-active nodes always stay in the *sleep* mode except for periodic energy and connectivity inspection instants. The frequency of those instants is variable and usually is much lower than that of the MAC layer beacons. In [15], nodes are assumed to be equipped with GPS devices, so that each node may periodically determine the number of nearby nodes within a pre-defined grid. Only one node within a local square may remain active, while the rest turn off their transceivers. In our scheme, the nodes also periodically turn off their transceivers to conserve energy. However, since each node determines whether or not to remain active solely based on the amount of its residual energy and the link-state information from the routing protocol, our scheme does not require geographical information.

Furthermore, in both [3] and [15], the network density (number of neighbors per node) under investigation is between 20 and 80. Therefore, the major performance gain of both approaches is materialized through reducing the very high network density. In contrast, the main goal of our scheme is to allow each node to *gradually* reduce its participation in networking activities as its stored energy is depleted. As demonstrated in Section VII, our approach can allow significant energy savings for networks of relatively low density (e.g., 10 neighbors per node).

III. RADIO TRANSCEIVER POWER CONSUMPTION

The proposed network participation criteria depend on the node's remaining energy level. Before describing these criteria, we first construct a model for ad hoc node energy consumption.

While the network radio transceiver is active, it consumes an average power of P_{active} , based on the proportion of the time that it spends in one of three active states. In the TRANSMIT state, a node is transmitting data at power $P_{transmit}$. Likewise, in the RECEIVE state, the node, while receiving a data packet, consumes power $P_{receive}$. Finally, in the IDLE state, the receiver consumes power P_{idle} , while sensing the channel for the arrival of new data. In order to conserve power, the transceiver may transition into SLEEP mode, consuming reduced power level of P_{sleep} .

Based on the results of previous research [12], we can make some valuable observations about the power consumption in each state. As one might expect, the transceiver consumes the most power while transmitting. However, transmitting *does not* consume considerably more power than receiving. Measurements from a WaveLAN PC card demonstrate only a 20% increase in power consumption between data transmission and data reception. It is also interesting to note that the transceiver's idle power consumption is only slightly less than the power consumption while receiving. Thus, the extra signal processing required for receiver demodulation and frame processing appears to be relatively inexpensive. In contrast, in sleep mode, the transceivers consume only about 10% of the active power levels.




Fig. 1. Two-state model for EDP

IV. ENERGY DEPENDENT PARTICIPATION

Since the transceivers consume a significant amount of energy while active, reducing the duration that the nodes are active can prolong their lifetime. This motivates our proposed Energy Dependent Participation criteria.

If a node wishes to maximize its battery life, the best policy is to leave its transceiver in the SLEEP mode or disconnect it altogether. However, such a policy is in conflict with the welfare of the whole network. Furthermore, it may work against the broader interests of the node itself, if any of its applications require services from the network.

In order to address the objectives of both the node and the network, we would like to have the node actively participate in the network at least some of the time. All else being equal, it seems reasonable that a node's network participation should be proportional to the amount of remaining energy (unless the node's network services are required by one of its applications or by the network itself; see Section V). The proportionality ensures that the batteries in all nodes are depleted at roughly the same rate over the long run. More precisely, let $E(t)$ be the amount of energy stored in a node at time t . Every τ seconds, the node decides to make its transceiver ACTIVE for the next τ seconds with probability $E(t)/E(0)$, or set it to SLEEP for the next τ seconds with probability $1 - E(t)/E(0)$.¹ Figure IV illustrates EDP represented by a two-state Markov chain.

In practice, the interval τ should be moderate, probably on the order of several seconds. If the interval is made extremely large (i.e., on the order of hours), then a node with fresh battery may end up with a continuously active transceiver, draining its battery too quickly. Short intervals are potentially a serious problem as well, as they encourage more frequent transitions between the ACTIVE and THE SLEEP modes. Each resulting change in network connectivity could trigger a flood of routing protocol updates, bogging down the network with control traffic and depleting the limited energy supply. We investigate in more detail the optimal value of τ in sections VII.

V. MAINTAINING NETWORK APPLICATIONS AND NETWORK CONNECTIVITY

The scheme described in the previous section attempts to gradually decrease a node's network participation in response to a reduced energy supply. While this may be a good strategy to extend battery life, its narrow goals can result in undesirable side-effects. For example, the node may inadvertently disconnect its

¹In this paper, we focus on homogeneous network applications, where all nodes are considered equal and fairness is desirable. In some heterogeneous networks, there may be practical considerations for regulating node participation based on the type of node. The models described in this paper can accommodate multiple classes of nodes by scaling the $E(t)/E(0)$ ratio appropriately.

own active networking application when it places its transceiver into the SLEEP mode. A transition into the SLEEP mode can also be damaging to the network, if the node’s absence leads to further network partitioning.

These issues can be addressed if the transceiver has access to the network layer and the routing protocol information. Network layer packets that have recently originated from or have been delivered to the given node indicate that it is an active endpoint in a data session. Link state information amassed by the routing protocol can be used to compute overall network connectivity and determine whether any node pairs rely on this node to relay communication. If the node determines that its participation is necessary, it can override a transition into sleep mode.

We devote the remainder of this section to the latter problem of detecting network partitioning due to node withdrawal. Specifically, a node X needs to determine whether there exists a node pair that is currently reachable, but would be disconnected as a result of X ’s withdrawal from the network. The “brute force” approach to solving this problem is to twice evaluate the connectivity between all node pairs²: first for the current network, and then for a network without X (and its adjacent edges). If the number of unreachable nodes increases after node X is removed, then node X is needed to maintain the current level of network connectivity.

We can significantly simplify this problem by recognizing that the withdrawal of node X from the network only eliminates the links that directly connect X with its neighbors. Consider a pair of X ’s neighbors (n_1, n_2) which are connected by the path $n_1 \rightarrow X \rightarrow n_2$. If n_1 can reach n_2 via a path that does not pass through X , then any pair of nodes connected by a path containing $n_1 \rightarrow X \rightarrow n_2$ are also connected by a path that does not pass through X . Since every path through node X contains a sub-path $n_i \rightarrow X \rightarrow n_j$, we can avoid further network partitioning if and only if all incoming neighbors of X can reach all outgoing neighbors of X without passing through X .

The above observation allows us to develop a computationally less complex potential partition checking algorithm. Assume that the network contains N nodes, each of which has an average of k neighbors. We begin by constructing an adjacency-list representation of the network graph, in $O(N \log N)$ time. For each incoming neighbor of X , we traverse the graph, at each hop verifying whether an outgoing neighbor of X has been visited. These checks have complexity $O(kNk) = O(Nk^2)$. The overall complexity is, therefore, $O(N \log N + Nk^2)$. For dense networks, this is computationally no worse than the “brute force” approach. However, many ad-hoc networks are relatively sparse, making this algorithm computationally more efficient.

If the network links are known to be bi-directional³, it is sufficient to simply check whether *one* of X ’s neighbors can reach

²For example, using all-pairs shortest paths algorithms such as Floyd-Warshall ($O(N^3)$) or Johnson ($O(N^2 \log N + N^2k)$), where N represents the number of network nodes and k the average number of neighbors per node

³In general, wireless connectivity may not be bi-directional, due to differences in transmitter power, receiver sensitivity, and channel interference. However, since link discovery and usage typically rely on a two-way handshake procedure, the “active” network links may all be bi-directional.




Fig. 2. Three-state model for EDP with partition checking

the remaining neighbors without going through X . If one neighbor can reach all others than the bi-directional links’ properties of commutativity and transitivity ensure that all pairs of neighbors are reachable. This reduces the complexity of the algorithm to $O(N \log N + Nk)$.

The node participation model described in Section IV consists of two basic states: SLEEP and ACTIVE. Two states are insufficient for partition checking, since a node in the SLEEP mode will not have an up-to-date view of the network topology. In order to avoid the problem of making decisions based on stale information, we require that after a SLEEP interval, nodes briefly WAKE-UP, in order to receive link state updates from its neighbors. After the WAKE-UP interval, the node decides between ACTIVE and SLEEP states, based on the remaining energy level, network application activity and potential partitioning. Figure 2 illustrates the additional WAKE-UP state that supports partition checking.

VI. EDP AND THE ROUTING PROTOCOL

The proposed EDP scheme is closely tied to the routing protocol. With partition checking, the optimal operation of an EDP enabled transceiver depends on both the energy consumption information and the network link-state information. Therefore, the network layer and routing protocol have direct impact on the performance of EDP.

The current ad hoc routing protocols can be roughly divided into three categories: *proactive*, *reactive*, and *hybrid*.⁴ A proactive routing protocol requires that each node maintains an up-to-date routing table, such that a route is readily available when data packets need to be send out.⁵ In reactive routing protocols, a node is not required to maintain a routing table (although route caches may be kept), but instead a route discovery process is initiated whenever it is needed.⁶ A hybrid routing protocol combines the benefit of both proactive and reactive approaches. For example, in the Zone Routing Protocol (ZRP) [6], a node proactively maintains the link-state information of only those nodes located within a surrounding variable-sized *routing zone*. This

⁴For a more in-depth discussion on the trade-off between various ad hoc routing protocols, the readers are referred to [6].

⁵Some examples of proactive routing protocols are DSDV, FSR, OLSR, STAR, TBRPF, and WRP.

⁶Some examples of reactive routing protocols are ABR, AODV, DSR, and TORA.

local information is leveraged to enable the *efficient* reactive discovery and maintenance of routes for destinations outside of its routing zone. This adaptive hybrid approach can significantly improve the scalability of ad hoc networks.

The effectiveness of EDP’s partition checking algorithm depends on the quality of the information provided by the routing protocol. Proactive link-state routing protocols provide each node with the current topology of the entire network, subject to the convergence of the routing protocol. For reactive routing protocols, each node maintains snapshots of some network links or paths in a route cache. If a path is in active use, the path’s links are likely to be valid at any time. On the other hand, links on inactive paths may no longer be valid⁷. In the case of hybrid routing protocols, nodes have access both to current local routing zone connectivity and snapshots of some distant network paths.

Hybrid routing’s limited view of local routing zones and active paths reduces the computational complexity of EDP partition checking. However, incomplete information may lead a node to incorrectly identify itself as the sole connector between two of its neighbors (i.e. the neighbors may be connected through an unknown path outside of the node’s local zone). In such situations, the node will unnecessarily remain active, wasting energy to keep its neighbors connected. Although a node may miss a legitimate SLEEP opportunity, as long as the available network information is valid, EDP will always catch and prevent partitioning caused by transitions to SLEEP mode. We show in Section VII that EDP performs well with a properly chosen local zone radius.

The actions resulting from the EDP algorithms impact the performance of the routing protocol by imposing higher demand on the network layer. In a proactive routing protocol, the sleeping and waking-up of the nodes can potentially incur more link updating traffic. In a reactive protocol, EDP can lead to more route querying traffic, when established routes break due to intermediate nodes entering the SLEEP mode. An increase in routing control traffic (including link updating packets and route querying packets) results in additional energy consumption, as well as potential network congestion. However, in a hybrid routing protocol such as ZRP, the effect of EDP can be moderated through the adaptive adjustment of the local zone radius. Our simulation results, presented in Section VII, show that the negative effect of EDP on the network layer is small when ZRP is employed.

VII. SIMULATIONS AND EVALUATIONS

A. Simulation Model

We have created an ad hoc network simulation environment to evaluate the proposed energy dependent ad hoc network participation scheme. Through simulations, we have investigated the effect of energy-inspection interval and partition-checking on battery lifetime, node survival rate, and network connectivity.

The simulated ad hoc network consists of 200 stationary nodes whose positions are chosen from a uniform random distribution over an area of 2 kilometers by 2 kilometers. Each node is

⁷Effective cache maintenance policies seek to minimize the presence of out-of-date route cache information.

state	power
$P_{transmit}$	0.285×10^{-6} J/bit, 1.425 J/sec
$P_{receive}$	0.925 J/sec
P_{idle}	0.925 J/sec
P_{sleep}	0.042 J/sec

TABLE I
POWER CONSUMPTION IN VARIOUS STATES

equipped with a network transceiver modeled after the Orinoco 802.11b PC Card.⁸ The transceivers operate at 11 Mbps and have an effective transmission radius of 250 meters⁹. Table I presents the numerical values of power consumption by a node at various states. These numbers correspond to the power ratio of 1 : 0.65 : 0.65 : 0.032 among transmitting (at full capacity), receiving, idling, and sleeping states, respectively.

A simplified version of the Zone Routing Protocol (without bordercasting [6]) is simulated at the network layer. A zone radius of 3 is used. Within each zone, the link-state update period is 10 seconds. For the transmission of each routing control packet, we account for the energy consumption of the transmitting node and all receiving *active* nodes within the transmission radius.

Data communication among the nodes is modeled as simultaneous network sessions. Every 50 seconds, 20 randomly chosen source-destination node pairs initiate independent sessions of duration 50 seconds and end-to-end bit rate 1.1 Mbps. In this preliminary study, we assume ideally shared channel media access. Thus, a session is *blocked* at the session initiation time if either the source or the destination node has no available battery energy, or if there is no route between the source-destination pair. After a session has begun, it is maintained as long as there is a route between the source and the destination. Namely, if the route carrying a session is broken due to an intermediate node going into the SLEEP mode, we assume that ZRP will provide an alternate route (if such a route exists) in time, such that the session may continue. Furthermore, the source and the destination nodes remain active during the session. Thus, a session is *dropped* over the session duration if either the source or the destination node has depleted its energy, or if all previously available routes are broken due to sleeping or dying intermediate nodes.

We simulate different scenarios where the sleep-cycle interval τ varies from 0.6 seconds to 600 seconds. At the beginning of each simulated scenario, each node has 8.4×10^3 joules of stored energy, which approximately equals to the energy of two fresh AAA batteries. We are primarily interested in the probability that a session survives (is neither blocked nor dropped) over its entire duration. Obviously, this probability decreases over time, as the network connectivity is reduced as each node’s battery is depleted.

B. Performance of EDP without Partition Checking

Figures 3-5 demonstrate the performance of the EDP scheme without partition checking.

⁸Additional simulations with a previous generation wireless network interface card (the 915 MHz WaveLAN Turbo 11 Mb PC Card) yield very similar results.

⁹This corresponds to roughly 10 neighbors per node.




Fig. 3. Average remaining energy, without partition checking




Fig. 4. Number of surviving nodes, without partition checking

Figure 3 illustrates how EDP can potentially improve the battery lifetime. When the transceiver is continuously in the ACTIVE state, as in the non-EDP case, the battery is completely drained after 2.3 hours. In contrast, the EDP enabled nodes *gradually* become less active network participants as the remaining energy reserves are diminished. This allows the battery lifetime to be more than doubled, in the case of $\tau = 60$ [sec], to 5 hours. Moreover, since the judicious usage of energy depends upon the accurate knowledge of the remaining energy, the battery lifetime can be increased by more frequent energy inspections.

Figure 4 further supports the above observation by illustrating the number of surviving nodes. In the non-EDP case, all nodes are lost after approximately 2.3 hours. However, with EDP, more nodes survive longer. For example, in the case of $\tau = 600$ [sec], 50% of the nodes survive after more than 2.9 hours, and in the case of $\tau = 60$ [sec], 50% of the nodes survive after 4.9 hours.

The above results show the effectiveness of EDP in energy preservation. However, being over-protective of the individual node's energy can lead to detrimental effects to the overall network performance. Figure 5 illustrates the trade-off between keeping nodes inactive and employing them to maintain the network connectivity. These plots show that, on the one hand, the session survival probability increases as τ is reduced from infinity (in the non-EDP case) to 600 seconds. On the other hand, the




Fig. 5. Session survival probability, without partition checking

session survival probability decreases as τ is further decreased from 600 seconds to 0.6 seconds. Therefore, although frequent energy inspections help prolong the battery lifetime, that does not necessarily translate into higher degree of network connectivity.

In particular, when the nodes switch between the active and sleeping states at a rate faster than the sessions are completed, the probability becomes significant that all routes between a source-destination pair are broken within a session duration. In those cases, the session dropping probability due to sleeping nodes in EDP can be much larger than that due to dying nodes in non-EDP. For example, in the case of $\tau = 0.6$ [sec], although the session survival probability is positive well beyond the 3-hour point, most of the time it is much inferior to the non-EDP case or the cases with larger values of τ . Section VII-C shows that the employment of partition checking can alleviate this problem.

C. EDP with Partition Checking

As explained in Section V, partition checking can improve the network connectivity in EDP. Figures 6-8 demonstrate the performance of EDP with partition checking.

Partition checking demands the active contribution of critical nodes even though their stored energy is low. Therefore, as shown in Figure 6, we see that the battery life is shortened, especially when the nodes frequently inspect their residual energy. For example, in the case $\tau = 60$ [sec], the battery life is 4.5 hours with partition checking, a 10% reduction from the non-partition-checking case. Figure 7 confirms that observation. It suggests that, in the case $\tau = 60$ [sec], only about 10 nodes survive at 4.9 hours, a 90% reduction from the non-partition-checking case.

However, the most important measure of the usability of a network is not the battery lifetime. Although partition checking leads to shorter battery lifetime, it provides significantly more robust network connectivity. Figure 8 illustrates that partition checking can vastly improve the performance of EDP when τ is not too large. For example, for the cases $\tau = 0.6$ [sec] with partition checking, at least 50% of the sessions survive up to 3.5 hours, which is about 33% longer than that best case (when $\tau = 600$ [sec]) without partition checking.

If we define the *usable lifetime* of an energy constrained network as the point where a session has 50% survival probability,




Fig. 6. Average remaining energy, with partition checking




Fig. 7. Number of surviving nodes, with partition checking

then, under the parameters considered in this work, using EDP with partition checking can result in up to 50% longer duration of usable lifetime compared with the traditional non-EDP ad hoc networks. The maximal performance gain is achieved when the sleep-cycle interval τ is approximately 0.6 seconds.

Figure 8 also suggests that further decreasing τ to below fraction of a second does not significantly increase the usable lifetime of the network. Indeed, if τ is extremely small, the resulting frequent link updating and route querying packets may overwhelm the network, leading to congestions and rapid energy depletion. The optimal value of τ largely depends on the system configuration.

VIII. CONCLUSIONS

Ad hoc network nodes are typically equipped with limited-capacity batteries. Considering this limited power supply, the wireless transceivers consume significant amount of energy in a relatively short period of time. Furthermore, it has been demonstrated that the energy consumption by a transceiver can be dominated by the amount of time that the transceiver is in the active but idling state. Therefore, the overall performance of an ad hoc network is closely related to the effectiveness of energy conservation at the transceivers.

In this work, we have proposed an Energy Dependent Participation (EDP) scheme that reduces the amount of time that the




Fig. 8. Session survival probability, with partition checking

transceivers are active. By allowing the transceivers to temporarily go into the sleep mode with probability proportional to the amount of consumed energy, EDP increases the battery lifetime and the node survivability. With partition checking, EDP further improves the network connectivity. Our simulation results have shown that EDP with partition checking based on the Zone Routing Protocol can double the the overall battery lifetime and increase the network usable lifetime (when the session survival rate drops below half) by over 50%.

REFERENCES

- [1] L. Bononi, L. Donatiello, "A distributed contention control mechanism for power saving in random access ad hoc wireless local area networks," *Proc. IEEE MoMuC*, November 1999.
- [2] J.-H. Chang and L. Tassiulas, "Energy conserving routing in wireless ad-hoc networks," *Proc. IEEE INFOCOM*, March 2000.
- [3] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris, "Span: an energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks," *Proc. ACM Mobicom*, July 2001.
- [4] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*, MIT Press, 1997.
- [5] J. Gomez, A. T. Campbell, M. Naghshineh, and C. Bisdikian, "PARO: Conserving Transmission Power in Wireless Ad Hoc Networks", *IEEE 9th International Conference on Network Protocols (ICNP'01)*, November 2001.
- [6] Z. J. Haas and M. R. Pearlman, "ZRP - a hybrid framework for routing in ad hoc networks," *Ad Hoc Networks*, Charles Perkins, ed., Addison Wesley Longman, 2001.
- [7] T. Imielinski, S. Viswanathan, and B. R. Badrinath, "Energy efficient indexing on air," *Proc. ACM SIGMOD*, pp. 25-36, May 1994.
- [8] A. Michail and A. Ephremides, "Energy efficient routing for connection-oriented traffic in ad-hoc wireless networks," *Proc. 11th IEEE PIMRC*, vol. 2, pp. 762-766, 2000.
- [9] V. Rodoplu and T. H. Meng, "Minimum energy mobile wireless networks," *J. Selected Areas Commun.*, vol. 17, no. 8, August 1999.
- [10] M. Sanchez, P. Manzoni, and Z. J. Haas, "Determination of critical transmission range in ad-hoc networks," *MMT'99*, October 1999.
- [11] S. Singh, M. Woo, and C. S. Raghavendra, "Power-aware routing in mobile ad hoc networks," *Proc. The Fourth ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom'98)*, October 1998.
- [12] M. Stemm and R. H. Katz, "Measuring and reducing energy consumption of network interfaces in hand-held devices," *IEICE Trans. on Commun. - Special Issue on Mobile Computing*, vol. E80-B, No. 8, 1997.
- [13] S. Tragoudas, "Power dissipation components of a management protocol for ad-hoc networks," *Proc. IEEE WCNC'00*, September 2000.
- [14] R. Wattenhofer, L. Li, P. Bahl, and Y.-M. Wang, "Distributed topology control for power efficient operation in multihop wireless ad hoc networks," *Proc. IEEE INFOCOM*, April 2001.
- [15] Y. Xu, J. Heidemann, and D. Estrin, "Geography-informed energy conservation for ad hoc routing," *Proc. ACM Mobicom*, 2001.