Power-Aware Routing in Mobile Ad Hoc Networks

Suresh Singh and Mike Woo
Department of ECE
Oregon State University
Corvallis, OR 97331
woo,singh@ece.orst.edu

C. S. Raghavendra
Aerospace Corporation
El Segundo, CA 90245
email: raghu@aero.org

Abstract

In this paper we present a case for using new *power-aware metrics* for determining routes in wireless ad hoc networks. We present five different metrics based on battery power consumption at nodes. We show that using these metrics in a shortest-cost routing algorithm reduces the cost/packet of routing packets by 5-30% over shortest-hop routing (this cost reduction is on top of a 40-70% reduction in energy consumption obtained by using PAMAS, our MAC layer protocol). Furthermore, using these new metrics ensures that the mean time to node failure is increased significantly. An interesting property of using shortest-cost routing is that packet delays do not increase. Finally, we note that our new metrics can be used in most traditional routing protocols for ad hoc networks.

1 Introduction

Ad Hoc networks are multi-hop wireless networks where all nodes cooperatively maintain network connectivity. These types of networks are useful in any situation where temporary network connectivity is needed, such as in disaster relief. An ad hoc network here would enable medics in the field to retrieve patient history from hospital databases (assuming that one or more of the nodes of the ad hoc network are connected to the Internet) or allow insurance companies to file claims from the field.

Building such ad hoc networks poses a significant technical challenge because of the many constraints imposed by the environment. Thus, the devices used in the field must be lightweight. Furthermore, since they are battery operated, they need to be energy conserving so that battery life is maximized. Several technologies are being developed to achieve these goals by targeting specific components of the computer and optimizing their energy consumption. For instance, low-power displays (see [13]), algorithms to reduce power consumption of disk drives (see [9, 19, 34]), low-power I/O devices such as cameras (see [5]), etc. all contribute to overall energy savings. Other related work includes the development of low-power CPUs (such as those used in laptops) and high-capacity batteries.

Our focus, in the past year, has been on developing strategies for reducing the energy consumption of the *communication subsystem* and *increasing the life* of the nodes. Recent studies have stressed the need for designing protocols to ensure longer battery life. Thus, [21] observes that the average life of batteries in an *idle* cellular phone is one day. [32] studies power consumption of several commercial radios (WaveLAN, Metricom and IR) and observes that even in Sleep mode the power consumption ranged between 150-170 mW while in Idle state the power consumption went up by one order of magnitude. In transmit mode the power consumption typically doubled. The DEC Roamabout radio [1] consumes approximately 5.76 watts during transmission, 2.88 watts during reception and 0.35 watts when idle.

If we examine the existing MAC protocols and routing protocols in this context we see a clear need for improvement: *in all of the current protocols, nodes are powered on most of the time even when they are doing no useful work.* At the MAC layer, nodes expend scarce energy when they overhear transmissions. In Figure 1, node A's transmission to node B is overheard by node C because C is a neighbor of A. Node C thus expends energy in receiving a packet that was not sent to it. In this case, clearly, node C needs to be powered off for the duration of the transmission in order to conserve its energy. Our MAC layer protocol (summarized in section 4) does precisely this and saves large amounts of energy. Routing protocols designed for ad hoc networks are also guilty of expending energy needlessly. In most of these protocols the paths are computed based on minimizing hop-count or delay. Thus, some nodes, become responsible for routing packets from many source–destination pairs. Over time, the energy reserves of these nodes will get depleted resulting in node failure. A better choice of routes is one where packets get routed through paths that may be longer but that pass through nodes that have plenty of energy reserves.

Our research has focussed on designing protocols that increase the life of nodes and the network. In order to produce a *complete solution*, we have attacked each layer (MAC,

MOBICOM 98 Dallas Texas USA
Copyright ACM 1998 1-58113-035-x/98/10...$5.00

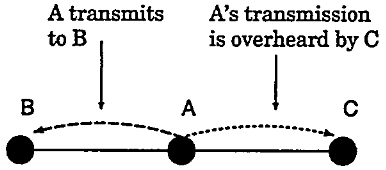

Figure 1: Unnecessary power consumption.

network and transport) individually. In our bottom-up approach, we optimize the energy consumption of the MAC layer first followed by the network layer and finally the transport layer. In [24] we present a MAC layer protocol for ad hoc networks that reduces energy consumption by 40% to 70% for different load and network conditions. An overview of this work is provided in section 4. **In this paper, we explore the issue of increasing node and network life by using power-aware metrics for routing.** Intuitively, it is best to route packets through nodes that have sufficient remaining power (rather than through a node whose battery is on its last legs). Similarly, routing packets through lightly-loaded nodes is also energy-conserving because the energy expended in contention is minimized. We show that power-aware routing (built on top of a power-aware MAC protocol) can save overall energy consumption in the network and, simultaneously, increase battery life at all nodes. Our work on optimizing transport layer protocols will be presented in an upcoming paper.

The remainder of this paper is organized as follows. In the next section we discuss the problem of routing in multi-hop wireless networks and provide a survey of metrics used by current routing protocols. In section 3 we discuss *different metrics* that result in power-aware routing. Section 4 outlines our energy conserving MAC layer protocol for multi-hop wireless networks. We also present related results on reducing energy consumption in cellular and wireless LAN environments by carefully designing the MAC protocol. Section 5 presents the results of our simulations where we demonstrate the use of new power-aware metrics. Finally, section 6 summarizes the main results and outlines our future research.

2 Metrics used in current Routing Protocols

The problem of routing in mobile ad hoc networks is difficult because of node mobility. Thus, we encounter two conflicting goals: on the one hand, in order to optimize routes, frequent topology updates are required, while on the other hand, frequent topology updates result in higher message overhead. Several authors have presented routing algorithms for these networks that attempt to optimize routes while attempting to keep message overhead small. In this section we breifly discuss the different *metrics* used for routing and then examine their effect on node and network life.

Different routing protocols use one or more of a small set of metrics to determine optimal paths. The most common metric used is *shortest-hop* routing as in DSR (Dynamic Source Routing [15]), DSDV (Destination Sequenced Distance Vector [26]), TORA (Temporally-Ordered Routing Algorithm [25]), WRP (Wireless Routing Protocol [22]) and in the DARPA packet radio protocol (see [16, 18]). Some of these protocols, however, can just as easily use *shortest delay* as the metric. *Link quality* is a metric that is used by SSA (Signal Stability based Adaptive Routing [8]) and by the DARPA protocol. Here, link quality information is used to select one among many different routes (in some cases a shortest-hop route may not be used because of poor link quality). In addition to link quality, SSA also uses *location stability* as a metric. This metric biases route selection toward routes with relatively stationary nodes. A benefit of these type of routes is that there will be little need to modify them frequently. Finally, the SRA protocol (Spine Routing Algorithm [7]) attempts to minimize the message and time overhead of computing routes. In this protocol, nodes are assigned to clusters (one or two-hops in diameter) and clusters are joined together by a virtual backbone. Packets destined for other clusters get routed via this backbone. The goal here is to reduce the complexity of maintaining routes in the face of node mobility. Of course, the routes are not necessarily the shortest.

The salient features of these protocols is summarized in Table 1. In this table, we have classified the protocols according to the metrics used for route optimization, the message overhead in determining routes, the type of protocol used and its convergence goals (*active* refers to a protocol that runs until all routing tables are consistent while *passive* refers to an algorithm that determines routes based on an as-needed basis).

2.1 Discussion of the power-awareness of current metrics

Some of these metrics, unfortunately, have a negative impact on node and network life by inadvertently overusing the energy resources of a small set of nodes in favor of others. For instance in the network illustrated in Figure 2, shortest-hop routing will route packets between 0–3, 1–4 and 2–5 via node 6, causing node 6 to die relatively early. Similarly, hierarchical and spine routing algorithms will (by their very design) exploit nodes that lie on the spine in order to reduce message overhead in routing table maintenance. In fact, it is important to observe that the metric of reducing message overhead may be misguided in the long-term. If we assume that 5-10% of network bandwidth is consumed by routing protocol overhead then reducing this number further will have little overall benefit if the data packets (that account for 90-95% of the bandwidth) either use sub-optimal routes or overextend the energy resources of a small set of nodes (on the spine, for instance). In fact, we can probably rephrase a version of Amdahl's Law (see pp. 29, [14]) for routing:

`Minimize the cost for the frequent case (data packets) over the infrequent case (control packets).`

Finally, we note that in most cases, link quality and location stability are orthogonal to the goal of power-awareness and therefore can be used *in conjunction* with the new metrics we define in the next section.

3 Metrics for Power-Aware Routing

Our key intuition in this paper is that conserving power and carefully sharing the cost of routing packets will ensure that node and network life are increased. However, we saw in the previous section that none of the metrics currently used

Protocol	Metrics	Message Overhead	Convergence	Protocol Type	Summary
DSR	Shortest Path	High	Passive	Source Routing	Route discovery, Snooping
DSDV	Shortest Path	High	Active	Distance Vector	Routing table exchange
DARPA	Shortest Path, Link Quality	High	Active	Distance Vector	Routing table exchange, Snooping
WRP	Shortest Path	High	Active	Distance Vector	Routing table exchanges
SSA	Location Stability, Link Quality	Moderate	Passive	Source Routing	Route Discovery
TORA	Shortest Path	Moderate	Passive	Link Reversal	Route update packets
SRA	Message and Time overhead	Moderate	Active	Hierarchical, Spine	Route discovery within cluster, Spine routing

Table 1: Comparison of several routing protocols for ad hoc networks.

for routing achieve this goal (in section 5 we support this claim via simulations). In this section, therefore, we present several power-aware metrics that do result in energy-efficient routes.

1. *Minimize Energy consumed/packet:* This is one of the most obvious metrics that reflects our intuition about conserving energy. Assume that some packet j traverses nodes $n_1, \ldots, n_k$ where n_1 is the source and n_k the destination. Let $T(a,b)$ denote the energy consumed in transmitting (and receiving) one packet over one hop from a to b. Then the energy consumed for packet j is,

   $$e_j = \sum_{i=1}^{k-1} T(n_i, n_{i+1})$$

   Thus, the goal of this metric is to,

   $$\text{Minimize } e_j, \forall \text{ packets } j \qquad (1)$$

   Discussion: It is easy to see that this metric will minimize the average energy consumed per packet. In fact it is interesting to observe that, under light loads, the routes selected when using this metric will be identical to routes selected by shortest-hop routing! This is not a surprising observation because, if we assume that $T(a,b) = T$ (a constant) , $\forall(a,b) \in E$, where E is the set of all edges, then the power consumed is $(k-1)T$. To minimize this value, we simply need to minimize k which is equivalent to finding the shortest-hop path.

   In some cases, however, the route selected when using this metric may differ from the route selected by shortest-hop routing. Thus, if one or more nodes on the shortest-hop path are heavily loaded, the amount of energy expended in transmitting one packet over one hop will not be a constant since we may expend variable amounts of energy (per hop) on contention. Thus, this metric will tend to route packets around congested areas (possibly increasing hop-count).

   One serious drawback of this metric is that nodes will tend to have widely differing energy consumption profiles resulting in early death for some nodes. Consider the network illustrated in Figure 2. Here, node 6 will be selected as the route for packets going from 0–3, 1–4 and 2–5. As a result node 6 will expend its battery resources at a faster rate than the other nodes in the network and will be the first to die. Thus, this metric does not really meet our goal of increasing node and network life.

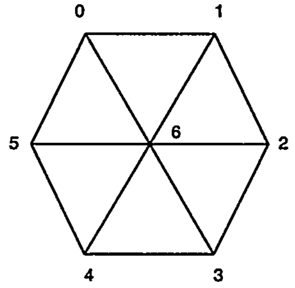

Figure 2: A network illustrating the problem with Energy/packet as a metric.

2. *Maximize Time to Network Partition:* This metric is very important in mission critical applications such as battlesite networks. Unfortunately, optimizing this metric is very difficult if we need to simultaneously maintain low delay and high throughput.

   Discussion: Given a network topology, using the max-flow-min-cut theorem, we can find a minimal set of nodes (the cut-set) the removal of which will cause the network to partition. The routes between these two partitions must go through one of these critical nodes. A routing procedure therefore must divide the work among these nodes to maximize the life of the network. This problem is similar to the "load balancing" problem where tasks need to be sent to one of the many servers available so that the response time is minimized – this is known to be an NP-complete problem. If we don't ensure that these nodes drain their power at equal rate, we will see delays increase as soon as one of these nodes die. Achieving equal power drain rate among these nodes require careful routing and is similar to the load balancing problem described above. In our case, since nodes in different partitions independently determine routes we cannot achieve the global balance required to maximize the network partition time while minimizing the average delay. We can also see that because the power consumption is dependent on the length of the packet we cannot decide optimal routes without the knowledge of future arrivals (similar to the knowledge of executing times of tasks in distributed systems). If all the packets are of same length, then we can ensure equal power drain rate among the critical nodes by selecting these nodes in a round-robin fashion in routing packets from one side to the other.

3. *Minimize Variance in node power levels:* The intuition behind this metric is that all nodes in the network are equally important and no one node must be penalized more than any of the others. This metric ensures that all the nodes in the network remain up and running together for as long as possible.

   Discussion: This problem is similar to "load sharing" in distributed systems where the objective is to minimize response time while keeping the amount of unfinished work in all nodes the same. Achieving this optimally is known to be intractable due to unknown execution times of future arrivals. Even if we are given a set of N tasks with variable lengths to be allocated to 3 or more machines, this problem is NP-complete as it is equivalent to the bin packing problem. A scheme that can be used to achieve the stated goal reasonably well is a policy called Join the Shortest Queue (JSQ). We can adopt such an idea by using a routing procedure where each node sends traffic through a neighbor with the least amount of data waiting to be transmitted. We can improve this further by doing some lookups of waiting traffic few hops away to decide the next best hop. An approximate routing procedure can be developed which uses the next hop based on total waiting traffic among its immediate neighbors when it has a choice. If all packets are of *same length*, however, then we can achieve this equal power drain rate by choosing next hop in a round-robin fashion so that on the average all nodes process equal number of packets.

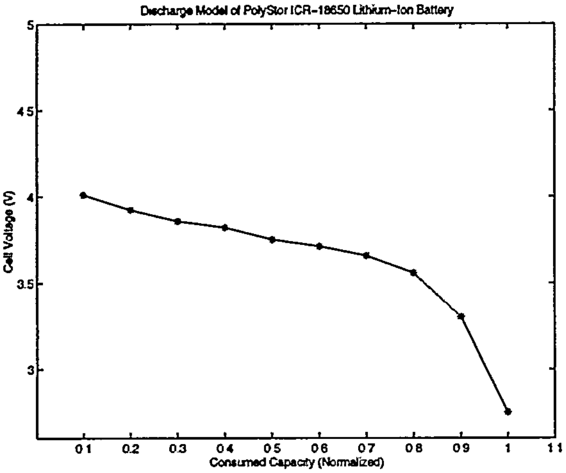

Figure 3: Example of a battery discharge function (Lithium-Ion).

4. *Minimize Cost/Packet:* If our goal is to maximize the life of all nodes in the network, then metrics other than energy consumed/packet need to be used. The paths selected when using these metrics should be such that nodes with depleted energy reserves do not lie on many paths. Let $f_i(x_i)$ be a function that denotes the *node cost or weight* of node i. x_i represents the total energy expended by node i thus far. We define the total cost of sending a packet along some path as the sum of the node weights of all nodes that lie along that path. The cost of sending a packet j from n_1 to n_k via intermediate nodes $n_2, \ldots, n_{k-1}$ is,

$$c_j = \sum_{i=1}^{k-1} f_i(x_i)$$

The goal of this metric is to,

$$\text{Minimize } c_j, \forall \text{ packets } j \tag{2}$$

Discussion: Intuitively, f_i denotes a *node's reluctance to forward packets* and we can see that with an appropriately chosen f_i, we can achieve different goals. Thus, if f_i is a monotone increasing function, then nodes (such as node 6 in Figure 2) will not be overused thus increasing their life. However, it is likely that the delay and the energy consumed/packet may be greater for some packets, such as those from 0–3, 1–4 and 2–5 that use 3-hop routes. This is not necessarily a drawback since the life of node 6 (in Figure 2) is increased and the variation in the lifetime of different nodes is reduced.

f_i can also be tailored to accurately reflect a battery's remaining lifetime. Many batteries display a discharge curve like the one illustrated in Figure 3 (see [12]). Here, we plot the normalized consumed capacity on the x-axis and the measured voltage on the y-axis. So, if the voltage is 2.8V, the battery is dead since all of its capacity (1 in normalized units) has been consumed. When the voltage is 3.6V, for example, 80% of the capacity has been consumed. One interesting choice for f_i is,

$$f_i(z_i) = \frac{1}{1 - g(z_i)}$$

where z_i denotes the measured voltage (that gives a good indication of the energy used thus far) and $0 \leq g(z_i) \leq 1.0$ is the normalized remaining lifetime (or capacity) of the battery ($(g(z_i), z_i)$ represents a point on the discharge curve). Using this type of a function ensures that the cost of forwarding packets is tied in closely with the power resources deployed in the network. Note that it is trivial to determine $f_i(z_i)$ since z_i can be read directly from the battery and the discharge curve is available for the battery[1].

An alternative form of f_i for this example (see Figure 3), however, is,

$$f_i(z_i) = \frac{1}{z_i - 2.8}$$

---

[1] We must add a word of caution though – in the case of older batteries, there is a significant error in determining the remaining lifetime from the voltage. This happens because of chemical degradation in the battery. One solution, for our purposes, would be to recompute the discharge curve as the battery ages or make available the discharge curves in some database that can be accessed by users based on their battery type, model and age.

this function has a reasonable node cost for about 80% of the battery's lifetime (the voltage drops from 4V to 3.6V) but after that point the cost grows rapidly. Intuitively, this form of f_i ensures that shortest-hop routing will be used when the network is new but as the network nodes near the end of their lifetimes, we carefully route packets so that no one node (or set of nodes) dies before the others (which can result in a partition).

Finally, we note that the discharge curve for some alkaline batteries is almost *linear* and we can associate a linear node cost function, such as,

$$f_i(z_i) = cz_i \tag{3}$$

with each node.

We can summarize some of the benefits of this metric as:

- It is possible to incorporate the battery characteristics directly into the routing protocol,
- As a side-effect, we increase time to network partition and reduce variation in node costs (though we do not optimize these metrics), and
- Effects of network congestion are incorporated into this metric (as an increase in node cost due to contention).

5. *Minimize Maximum Node Cost:* Let $C_i(t)$ denote the cost of routing a packet through node i at time t. Define $\hat{C}(t)$ denote the maximum of the $C_i(t)$s. Then,

$$\text{Minimize } \hat{C}(t), \forall t > 0 \tag{4}$$

metric minimizes maximum node cost. An alternative definition is to minimize the maximum node cost *after routing N packets to their destinations or after T seconds.* All of these variations ensure that node failure is delayed and a side effect is that the variance in node power levels is also reduced. Unfortunately, we see no way of implementing this metric directly in a routing protocol but minimizing cost/node does significantly reduce the maximum node cost (and hence time to first node failure).

The five metrics discussed above do, in different ways, express our intuition about conserving energy in the network by selecting routes carefully. However, what *protocols* best implement these metrics? It is easy to see that any protocol that finds shortest paths can be used to determine optimal routes based on the first and fourth metrics discussed above (equations 1, 2). To implement the first metric, we simply associate an edge weight with each edge in the network. This weight reflects the value $T(a,b)$. For the second metric (cost/packet), we associate node weights f_i with each node and compute the shortest path as usual. We have not yet implemented the other three metrics but we have determined that they are optimized somewhat by the metric (cost/packet) if we select f_i's carefully.

Finally, it is important to point out that our metrics do not necessarily need to be used for routing all the time. Rather, when the network is new (when all nodes are replete with energy resources), shortest-hop routing can be used. However, after some time when energy resources have fallen below a threshold, nodes can begin using one of the above routing metrics. Another related point is that routing protocols might use these metrics for routing most packets but switch to shortest-hop (or delay) routing for a fraction of the packets that have a high priority.

4 Overview of PAMAS (*Power-Aware Multiple Access protocol with Signalling*)

In this section we provide an overview of our MAC layer protocol for ad hoc networks. We use this protocol as the MAC protocol in our simulator as well. **Thus, the energy savings reported in section 5 are savings that are obtained *on top of* the considerable savings due to PAMAS.** The PAMAS protocol saves 40-70% of battery power by intelligently turning off radios when they cannot transmit or cannot receive packets. Thus, in the scenario illustrated in Figure 1, node C powers itself off for the duration of the transmission from A to B. Node C will thus conserve its battery power because it will not expend energy in listening to A's transmission. The specific conditions under with nodes power off in PAMAS are:

- A node powers off if it is overhearing a transmission and does not have a packet to transmit,
- If at least one neighbor is transmitting and at least one neighbor is receiving a transmission, a node may power off. This is because, even if the node has a packet to transmit, it cannot do so for fear of interfering with its neighbor's reception,
- If all of a node's neighbors neighbors are transmitting (and the node is not a receiver), it powers itself off.

A fundamental problem that arises when nodes power themselves off is, *for how long can a node remain powered off?* In the optimal case, a node powers itself off exactly when one of the conditions above holds true. However, in actual implementation, a node needs to estimate this length of time (keep in mind that a node cannot sense carrier when it is powered off so it has no way of knowing when a transmission in its neighboorhood has completed). In our protocol, as in all other MAC layer protocols for ad hoc networks, nodes attempt to grab the channel by exchanging RTS/CTS (ready to send and clear to send) messages. Thus, the sender transmits a RTS message. The receiver responds with a CTS message if it received the RTS message uncorrupted. The sender begins transmission upon receiving the CTS. In PAMAS, this exchange of RTS/CTS messages takes place over a separate signalling channel[2]. Thus, this exchange does not affect any ongoing data transmissions. The RTS/CTS messages contain the length of the packet the sender will send. Thus, any other node in the neighborhood can determine the length of the transmission and power off if one of the above conditions is met. A problem arises in the case when a node that has powered itself off wakens to hear a new ongoing transmission. In this case, it needs to be able to estimate the length of the remaining transmission and

---

[2] In PAMAS the receiver transmits a busy tone once it begins hearing the packet. This is done to combat a specific hidden-terminal problem.

power itself off (if one of the conditions above is met) again. We have a protocol that runs over the signalling channel that allows nodes to query transmitters about the length of the remaining transmission. Collisions during this enquiry (which are likely in high-degree networks since several nodes may power off as a consequence of a transmission and may waken simultaneously) are handled with a modified binary backoff algorithm. This algorithm can be tuned so that overhead of the algorithm is traded off against accuracy in the estimate of the length of the remaining transmission.

Figure 4 illustrates the power savings obtained (as a percentage) when using PAMAS. The network used is a 20-node random network. The x-axis denotes the edge probability. Different curves indicate power savings for different network loads. Note that at high loads the power savings are smaller because a large amount of power is consumed in contention. The savings, however, increase with increasing node connectivity since a node has more opportuinities to power-off. The PAMAS protocol is non-trivial and we cannot explain its operation in any detail here. However we would like to point out that in PAMAS the *delay* and *throughput* are not changed even when nodes power off. This is because the conditions under which nodes power off are such that the node powering off cannot transmit or receive packets anyway. A detailed discussion of PAMAS is provided in [23]. We have derived bounds on the maximum achievable power savings in [24].

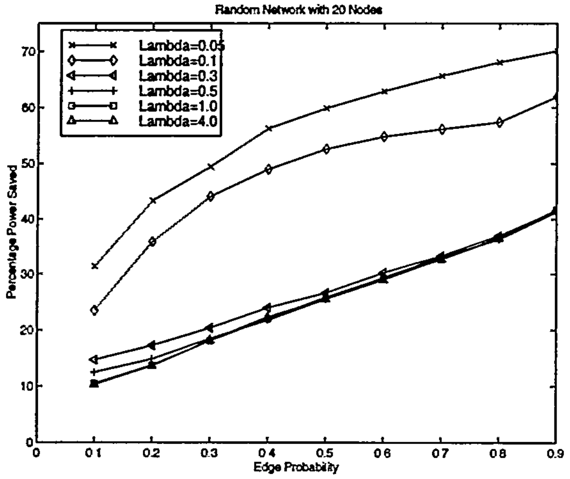

Figure 4: Power saved in random networks with 20 nodes.

4.1 Related Work on Power-Conserving MAC Protocols

Recently, some researchers have begun studying the problem of reducing power consumption by the wireless interface in **single-hop** wireless networks. Most approaches are based on the paging protocols POCSAG and FLEX where a base station periodically transmits a beacon followed by a minislot containing the ID of nodes that have a page waiting for them. These nodes remain awake in order to receive their messages while all the others power themselves off. A similar idea (based on reservation) is included in the IEEE 802.11 standard as well (see [29]). Here, nodes transmit their requests to the base station during specific reservation intervals and the base station transmits a TIM (Traffic Indication Map) that includes the transmission schedule for the nodes. All nodes not participating in transmission or reception of packets go into doze mode until the next reservation period. The standard also includes an extension of this idea to ad hoc single-hop networks. Here, nodes compete to be elected the leader to play the role of the base station. [30] presents a comparison of the power consumption behavior of three protocols - IEEE 802.11, DQRUMA (see [20]) and DSA++ (see [27]) - in a single-hop environment. Their main conclusions are that contention results in higher energy consumption while reservation and scheduling results in lower energy consumption. [6] also discusses the energy consumption of protocols and shows that persistence is not always a good choice and adaptive strategies that avoid packet retransmissions during bad channel periods is a good energy conserving strategy. Furthermore, [6] presents a access protocol for cellular networks based on ALOHA and reservation (the protocol is similar to IEEE 802.11) and analyze its performance (energy consumed and throughput). [31] also presents a reservation-based power conserving access protocol for mobile ATM networks.

5 Validation of the Power-Aware Metrics

We conducted extensive simulations to better understand the properties of the new metrics and the effect of using these metrics on end-to-end packet *delay*. Specifically, we compared the performance of ***shortest-hop*** routing with ***shortest-cost*** routing (equation 2) and quantified the difference between these two approaches using three measures[3]:

1. End-to-end packet delays (measured as the difference between time when a packet enters the system and time when it finally departs),
2. Average cost/packet (measured for each packet), and
3. Average maximum node cost (computed after 300 seconds of simulation time)

For the shortest-cost routing approach, we used several different f_i functions. In this paper, however, we only present two of these models for f_i. The first model was a *linear* model where $f(x) = cx$ for some constant $c < 1$ and the second model was a *quadratic* model where $f(x) = cx^2$. The linear model is based on the discharge curve of alkaline batteries while the quadratic model represents the precipitous discharge in battery life for lithium-ion batteries (Figure 3).

For the simulation, we used a 16-node mesh topology and 10 and 20-node random graphs. The random graphs were generated as follows. For each pair of possible edges, we toss a coin that has a probability p of coming up heads. If it does come up heads, we put that edge in otherwise we leave it out. We varied the value of p from 0.1 to 0.5. Intuitively, $p = 0.1$ produces a sparse graph while $p = 0.5$ produces a dense graph. We only considered connected networks in this study and we did not include node mobility. The reason we did not account for mobility is because we were not actually simulating a routing protocol (whose performance would depend on the mobility model) but only evaluating different power-aware metrics.

---

[3] We did not consider hierarchical spine routing because of our criticism in section 2

Packets arrive at each node according to a poisson process. The packet arrival rate λ varies between 0.05 and 0.5 packets/sec/node. Each node maintains a FIFO buffer of packets that need to be forwarded to the next hop. Every packet is timestamped when it first enters the system and then again when it arrives at its destination allowing us to compute delays. Further, node costs are updated constantly and when a packet is transmitted over one hop, we add the current node cost to the total cost of the packet. The packet costs are averaged out at the end of the simulation as are the node costs.

We ran each simulation 20 times and computed the mean and the standard deviation for each of the three metrics mentioned earlier (delay, cost/packet and average max node cost) for shortest-hop routing and shortest-cost routing. In the graphs we plot the *percentage improvement* in these metrics when we use shortest-cost routing. We have not plotted the curves for delay because there was *no difference in the average packet delay* (computed separately for packets travelling over one hop, two hops, etc.) between shortest-hop routing and shortest-cost routing. This result was surprising because we had expected a slight worsening in delay for packets (in the shortest-cost case) as they get routed around nodes with high cost (or low remaining lifetime). On closer examination of the simulation trace we found that some packets did indeed take longer routes and of these some did have higher delay (measured in time steps). However, the number of these packets was not large and as a result did not contribute to a statistically significant result. What was more significant, under high loads, was the fact that shortest-hop routing resulted in slightly longer packet delays (because of congestion) while shortest-cost routing (which is a function of energy consumed and is hence affected by contention costs) resulted in shorter delays since congested routes were not chosen! So, overall, we conclude that *packet delay* is unaffected when using shortest-cost routing.

Let us now consider the relative improvement in the *cost/packet* and *max node cost* metrics when using shortest-cost routing. We need to mention that both the shortest-hop and shortest-cost simulations were run on top of PAMAS. Thus, the improvement we see is in addition to the improvement gained by PAMAS (which is significant). Let us first look at a 10-node random network. Figure 5 illustrates the percentage improvement in the cost/packet/hop for different values of p. Each curve represents a different value of λ. The plot on the left shows the improvement when we use a linear cost function for f and the plot on the right shows the improvement when the cost function is quadratic. We can see that the improvement is in the 5-15% range. Figure 6 illustrates the same set of plots for 20-node random networks.

It is interesting to observe that the *savings are greater in larger networks*. This is not surprising because larger networks have more routes to choose from. A second observation we can make is that *savings increase with load*. This is because at very low loads, the cost differential between nodes is too small to matter. However as load increases, this cost differential increases and is reflected in cost savings per packet. Interestingly however, at heavy loads (beyond 0.2 or 0.3 in these studies), the improvement remains constant and, in fact, becomes negligible at very high loads (overloaded conditions). This last graph (with $\lambda = 1.5$ packets/node/sec) was not plotted because the savings were zero. The reason for this is that all nodes have a full buffer and expend huge amounts of energy in contention which results in reducing the node cost differential. Finally, we observe that the savings in cost *increases with edge probability p*. The reason for this is that at small p, the network is sparse resulting in few alternative routing paths while at higher p, more paths become available. The cost function f also affects the savings in cost. As the graphs show, savings are greater for the quadratic cost function than for the linear. This is because the cost differential between nodes increases sharply with a quadratic function.

We plot the reduction in maximum node costs for 10-node and 20-node random networks in Figures 7 and 8. In the 10-node network, there is a 5-10% reduction in maximum node cost for the linear case and 5-50% for the quadratic case. These numbers become 5-45% for the linear case and 15-120% for the quadratic case when we have a 20-node network. The reasons for this dramatic increase in savings in larger networks is because of the availability of more routes. Likewise, the savings increase in denser networks and they increase (initially) with λ. All for the same reasons as discussed previously.

Figure 9 illustrates the cost savings per packet and the reduction in maximum node cost for a 16-node mesh. We used the mesh because it provides with a well-connected topology and allows us to verify our conclusions from the random network topologies. As we can see, as the load increases (along the x-axis), the savings in cost per packet increase at first and then decreases as load continues to increase. The reason for the initial increase is that at very low loads, node costs are almost the same. As load increases, there is an increasing difference in node costs between shortest-hop and shortest-cost routing. Finally, at very high loads, the cost of all nodes is almost the same and thus there are no savings. The same behavior is illustrated in the plot on the right where we show the reduction in maximum node cost.

5.1 Summary of Results

Based on the simulations, we can conclude that using power-aware metrics to find routes is very beneficial because the difference in battery consumption between various nodes is reduced. This typically means longer network life and longer time to node failure. The specific conclusions from the experiments are:

1. Larger networks have higher cost savings,
2. Cost savings are best at moderate network loads and negligible at very low or at very high loads,
3. Denser networks exhibit more cost savings in general, and
4. The cost function used dramatically affects the amount of cost savings.

It is worth pointing out that our results will hold true in networks where nodes are mobile. This is because nodes in real networks do not move randomly independently. Rather, clusters of nodes move in correlated ways (image a platoon of soldiers). If, however, nodes do move randomly independently, then we believe that there will be small, if any, cost savings obtainable by using power-aware metrics (note, however, that PAMAS will still deliver huge savings).

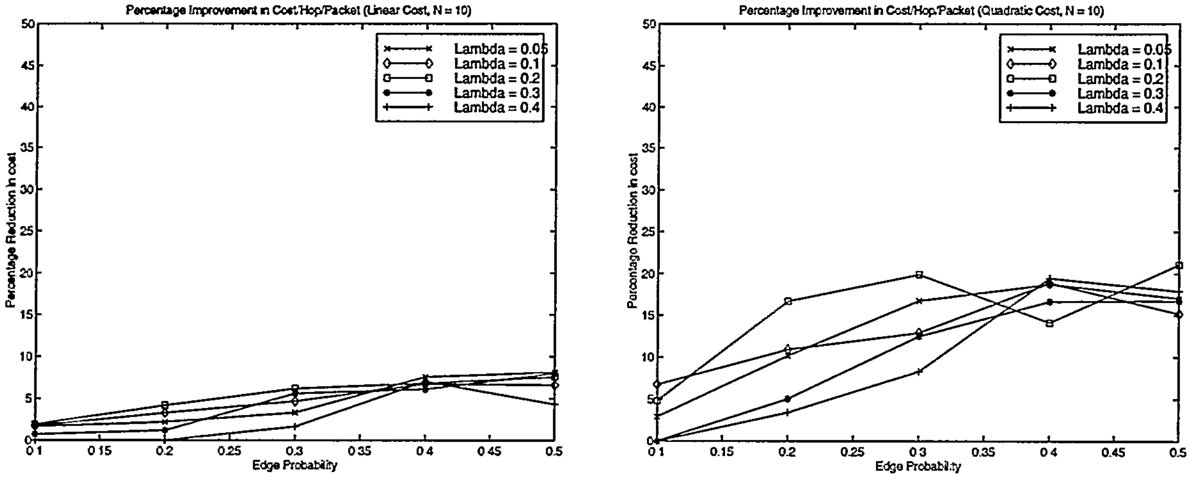

Figure 5: Percentage reduction in average cost in 10-node random networks.

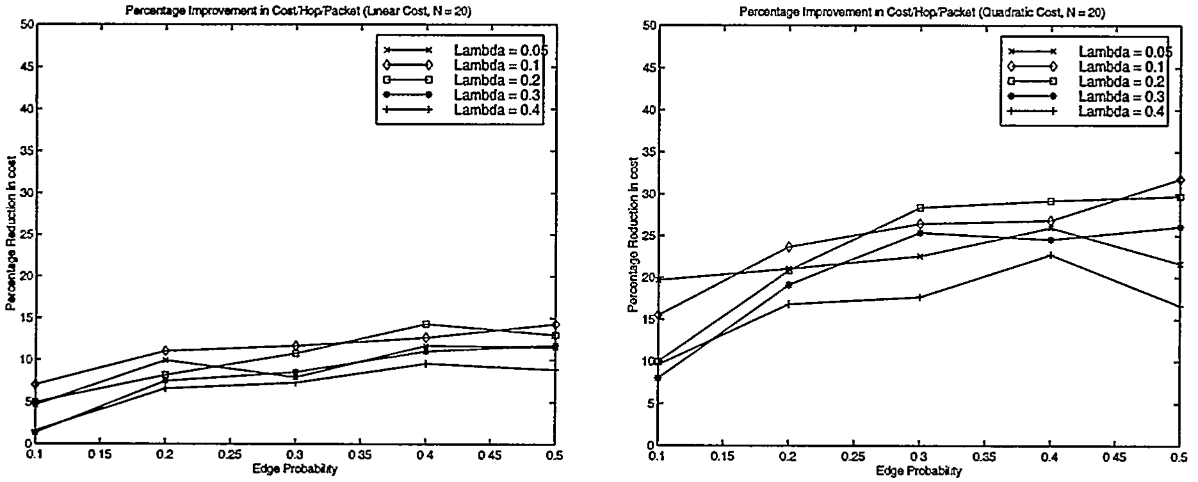

Figure 6: Percentage reduction in average cost in 20-node random networks.

6 Conclusions

In this paper we discussed the need to make routing protocols power-aware. Thus, rather than using traditional metrics such as hop-count or delay for finding routes, we believe that is more important to use cost/packet and maximum node cost (which are functions of remaining battery power) as metrics. Our simulations demonstrated that significant reductions in cost can be obtained by using shortest-cost routing as opposed to shortest-hop routing. A feature of our metrics is that they can be easily incorporated for use in existing routing protocols for ad hoc networks.

Acknowledgements

This research was supported by NSF grant NCR-9706080.

References

[1] `http://www.networks.digital.com/npb/html/products_guide/roamwir2.html`, Jan 14, 1998.

[2] D. Beyer, "Accomplishments of the DARPA SURAN Program", *Proc. IEEE MILCOM'90*, Monterey, CA, Oct. 1990.

[3] V. Bharghavan, A. Demers, S. Shenkar and L. Zhang, "MACAW: A Media Access Protocol for Wireless LANs", *Proceedings ACM SIGCOMM'94*, pp. 212-225, 1994.

[4] Béla Bollobás, Random Graphs, Academic Press, 1985.

[5] A. Chandrakasan, T. Simon, J. Goodman and W. Rabiner, "Signal Processing for an ultra low power Wireless Video Camera", *3rd International Workshop on Mobile Multimedia Communications*, Princeton, NJ, September 25-27, 1996.

[6] A. Chockalingam and M. Zorzi, "Energy Consumption Performance of a class of Access Protocols for Mobile Data Networks", *Proc. IEEE VTC'98*, Ottawa, Canada, May 18-21, 1998.

[7] B. Das, R. Sivakumar and V. Bharghavan, "Routing in Ad Hoc Networks using a Spine", *Proc. IEEE ICCCN'97*, (1997).

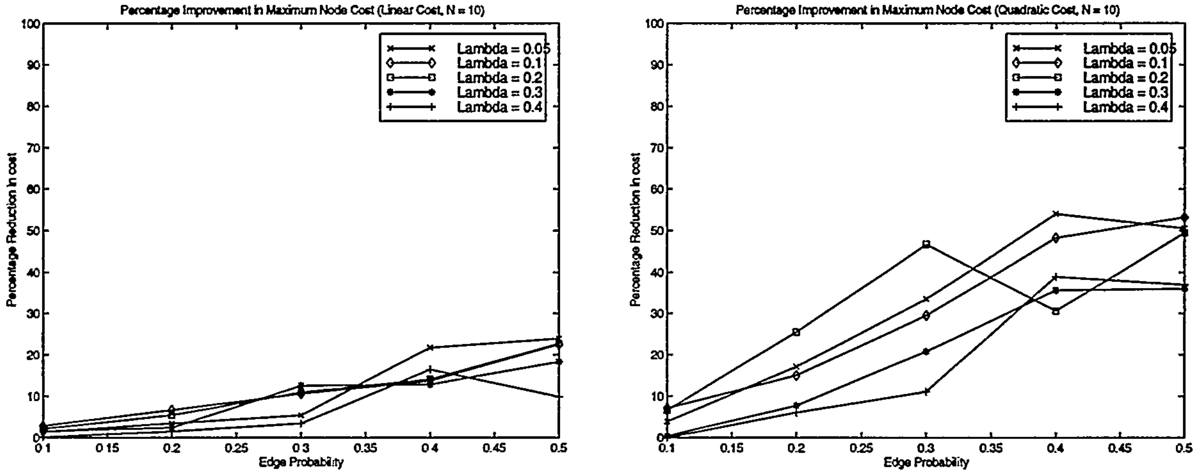

Figure 7: Percentage reduction in maximum node cost in 10-node random networks.

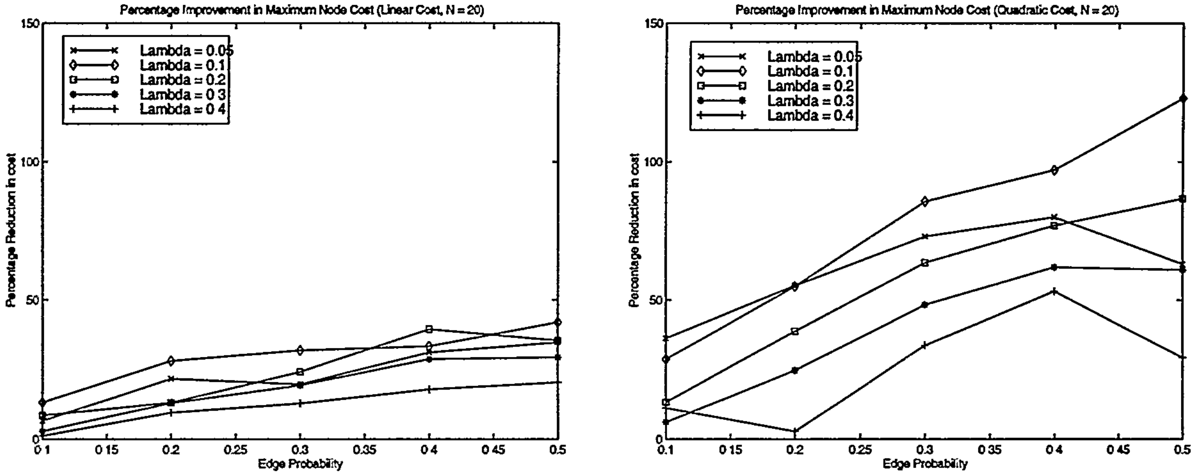

Figure 8: Percentage reduction in maximum node cost in 20-node random networks.

[8] R. Dube, C. D. Rais, K-Y Wang and S. K. Tripathi, "Signal Stability-based Adaptive Routing (SSA) for Ad Hoc Mobile Networks", Technical Report CS-TR-3646, UMIACS-TR-96-34, August 28, 1996.

[9] F. Douglis, F. Kaashoek, B. Marsh, R. Caceres, K. Lai and J. Tauber, "Storage Alternatives for Mobile Computers", *Proc. 1994 Symposium on Operating Systems Design and Implementation*, OSDI, November 1994.

[10] Chane L. Fullmer and J.J. Garcia-Luna-Aceves,"Solutions to Hidden Terminal Problems in Wireless Networks", *Proceedings ACM SIGCOMM'97*, Cannes, France, Sept. 14-18, 1997.

[11] Mario Gerla and J.T.-C. Tsai,"Multicluster, Mobile, Multimedia Radio Network", *ACM-Baltzer Journal of Wireless Networks*, Vol. 1, No. 3, pp. 255-265, 1995.

[12] S. Gold,"A PSPICE Macromodel for Lithium-Ion Batteries", *The 12th Annual Battery Conference on Applications and Advances*, California State Univ., Long Beach, CA, Jan 14-17, 1997. (see Fig. 4)

[13] E.P. Harris and K.W. Warren,"Low Power Technologies: A System Perspective", *3rd International Workshop on Mobile Multimedia Communications*, Princeton, NJ, September 25-27, 1996.

[14] J. L. Hennessy and D. A. Patterson,Computer Architecture A Quantitative Approach, 2nd edition, Morgan Kaufmann, San Francisco, CA, (1996).

[15] David B. Johnson and David A. Maltz. "Dynamic source routing in ad hoc wireless networks", In Tomasz Imielinski and Henry F. Korth, editors, Mobile Computing, pages 153-181. Kluwer Academic Publishing, 1996.

[16] John Jubin and Janet D. Tornow, "The DARPA packet radio network protocols", *Proceedings of the IEEE*, 75(1):21-32, January 1987.

[17] P. Karn,"MACA - a New Channel Access Method for Packet Radio", in *ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, pp. 134-140, 1990.

[18] Gregory S. Lauer. "Packet-radio routing", In Martha Steenstrup, editor, Routing in Communications Networks, Chapter 11, pages 351-396. Prentice-Hall, 1995.

[19] K. Li, R. Kumpf, P. Horton and T. Anderson,"A Quantitative Analysis of Disk Drive Power Management in Portable

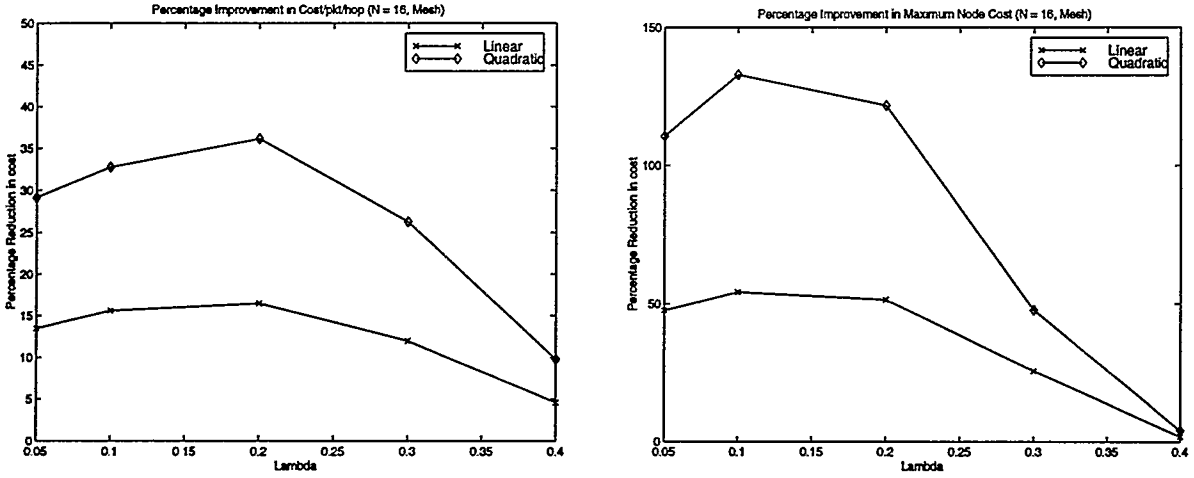

Figure 9: Percentage reduction in cost/pkt/hop and maximum node cost in a 16-node mesh.

Computers", *Proceedings 1994 USENIX*, San Francisco, CA, pp. 279-291, 1994.

[20] M. J. Karol, Z.Liu and K. Y. Eng, "Distributed-queuing request update multiple access (DQRUMA) for wireless packet (ATM) networks, *Proc. IEEE ICC'95*, June 1995, pp. 1224-1231.

[21] W. Mangione-Smith and P.S. Ghang, "A low power medium access control protocol for portable multi-media systems", *3rd International Workshop on Mobile Multimedia Communications*, September 25-27 1996.

[22] S. Murthy and J.J. Garcia-Luna-Aceves,"An Efficient Routing Protocol for Wireless Networks", *ACM Mobile Networks and Applications Journal*, Special Issue on Routing in Mobile Communications Networks, 1996.

[23] S. Singh and C. S. Raghavendra,"Power efficient MAC protocol for multihop radio networks", *Proc. IEEE PIMRC'98*, Sept, 1998.

[24] S. Singh and C. S. Raghavendra,"PAMAS – Power Aware Multi-Access protocol with Signalling for Ad Hoc Networks", *ACM Computer Communication Review*, July 1998.

[25] V. D. Park and M. S. Corson,"A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks", *Proc. IEEE INFOCOM'97*, Kobe, Japan, (1997).

[26] Charles E. Perkins and Pravin Bhagwat, "Routing over multi-hop wireless network of mobile computers", In Tomasz Imielinski and Henry F. Korth, editors, Mobile Computing, pages 183–205. Kluwer Academic Publishing, 1996.

[27] D. Petras and A. Krämling, "MAC protocol with polling and fast collision resolution for an ATM air interface", *Proc. IEEE ATM'96 workshop*, San Francisco, CA, Aug. 1996.

[28] Chunhung Richard Lin and Mario Gerla,"Asynchronous Multimedia Multihop Wireless Networks", *Proceedings IEEE INFOCOM'97*, (1997).

[29] Christian Röhl, H. Woesner and A. Wolisz, "A Short Look on Power Saving Mechanisms in the Wireless LAN Standard Draft IEEE 802.11", *Proc. of the 6th WINLAB Workshop on Third Generation Wireless Systems*, March 1997.

[30] Krishna M. Sivalingam, M. B. Srivastava and P. Agrawal,"Low Power Link and Access Protocols for Wireless Multimedia Networks", *Proc. IEEE Vehicular Technology Conference VTC'97*, Phoenix, AZ, May 4-7, 1997.

[31] Krishna M. Sivalingam, M. B. Srivastava, P. Agrawal and J-C. Chen, "Low-Power Access Protocols Based on Scheduling for Wireless and Mobile ATM Networks", *Manuscript*, http://www.eecs.wsu.edu/~krishna.

[32] M. Stemm and P. Gauthier and D. Harada,"Reducing power consumption of network interfaces in hand-held devices",*3rd International Workshop on Mobile Multimedia Communications*,September 25-27, 1996.

[33] C-K Toh,"The Cambridge Ad Hoc Mobile Routing Protocol", in Wireless ATM and Ad Hoc Networks, Kluwer Academic Publishers, Boston, Chapter 9, (1997).

[34] S. Zdonik, M. Franklin, R. Alonso and S. Acharya,"Are "disks in the air" just pie in the sky?", *IEEE Workshop on Mobile Computing Systems and Applications*, Santa Cruz, CA, pp. 12-19, December 1994.

[35] Michele Zorzi and R. R. Rao, "energy Management in wireless Communications", *Proc. 6th WINLAB Workshop on Third Generation Wireless Information Networks*, March 1997.