

Exploiting Heterogeneity in Sensor Networks

Mark Yarvis[†], Nandakishore Kushalnagar[†], Harkirat Singh^{✧†}, Anand Rangarajan[†], York Liu[†], Suresh Singh[✧]

[†] Intel Research and Development
2111 N.E. 25th Ave., Hillsboro, OR 97124

[✧] Computer Science Dept., Portland State University
1900 SW 4th Ave., Portland, OR 97201

Abstract—The presence of heterogeneous nodes (i.e., nodes with an enhanced energy capacity or communication capability) in a sensor network is known to increase network reliability and lifetime. However, questions of *where, how many, and what types* of heterogeneous resources to deploy remain largely unexplored. We focus on energy and link heterogeneity in ad hoc sensor networks and consider resource-aware MAC and routing protocols to utilize those resources. Using analysis, simulation, and real testbed measurements, we evaluate the impact of number and placement of heterogeneous resources on performance in networks of different sizes and densities. While we prove that optimal deployment is very hard in general, we also show that only a modest number of reliable, long-range *backhaul* links and *line-powered* nodes are required to have a significant impact. Properly deployed, heterogeneity can triple the average delivery rate and provide a 5-fold increase in the lifetime (respectively) of a large battery-powered network of simple sensors.

I. INTRODUCTION

Early research in ad hoc wireless sensor networks focused on networks in which all nodes possess identical software and hardware. This homogeneous architecture is attractive because it is resilient to individual failures. More recently, however, heterogeneous sensor networks have become popular, particularly in real deployments because of their potential to increase network lifetime and reliability without significantly increasing the cost.

For example, Intel has deployed a pilot application of sensor networks to monitor the health of mechanical equipment in its fabrication plants [7]. While sensors attached to pumps and motors in the fabrication plant may be line powered, use of battery-power sensors reduces installation cost and complexity. These sensors must conserve power, minimizing computation and communication. Additional nodes placed where standard wall sockets are available can incorporate high-speed microprocessors and high-bandwidth, long-distance network transceivers. These nodes can provide in-network data processing, longer-term storage, and access to building network infrastructures, such as Ethernet or 802.11 networks. The opportunities for heterogeneity in this example are typical of indoor sensing applications and also common in a wide variety of deployment scenarios.

In heterogeneous sensor networks, typically, a large number of inexpensive nodes perform sensing, while a few expensive nodes (perhaps embedded PCs) provide data filtering, fusion and transport. This partitioning of tasks ensures a cost-effective design as well as a more efficient implementation of the overall sensing application. However, realizing the full

potential of heterogeneity requires careful network engineering including careful placement of the heterogeneous resources and the design of resource-aware protocols. Effective exploitation of heterogeneity can impact the economic feasibility of applications such as industrial equipment monitoring.

We identify three common types of hardware heterogeneity: *computational heterogeneity* where some nodes have added computational power (e.g., Intel's Stargate), *link heterogeneity* where some nodes have long-distance highly reliable communication links (e.g., 802.11 connectivity), and *energy heterogeneity* where nodes have unlimited energy resources (e.g., connection to a wall outlet). In this paper we focus on the largely unexplored questions of *where, how many, and what types* of heterogeneous resources to deploy to maximize benefit.

We present analytical, simulation, and real testbed results to evaluate the impact of energy heterogeneity and link heterogeneity on sensor networks, which are typically designed to gather data from all nodes to a single sink. We evaluate analytically the theoretical benefit of link and energy heterogeneity. We then prove that optimal placement of resources in an arbitrary network is computationally expensive (NP-hard in some cases). We develop techniques for utilizing heterogeneity, such as (1) effective placement strategies, (2) MAC-level support to exploit the benefits of these resources, and (3) resource-aware routing. Finally, through testbed and simulation experiments, we evaluate the impact of the degree and placement of heterogeneous resources in networks of various sizes and densities. Our results indicate that a modest number of heterogeneous resources can triple the average delivery rate and provide a five-fold increase in the lifetime (respectively) of a large battery-powered network of simple sensors.

II. RELATED WORK

Use of heterogeneity in ad hoc networking is not new. Researchers at Millennial Net have developed a hardware and software architecture that presumes the presence of heterogeneous energy sources and communication capabilities [14]. References [10] and [16] provide examples of two real sensor network deployments that utilize heterogeneous nodes for processing and transport tasks. Researchers have demonstrated the necessity of heterogeneity, deployed heterogeneous networks, and described mechanisms that bias the use of some nodes over others for packet forwarding and processing. However, to our knowledge there has been no comprehensive study evaluating strategies for effective use of heterogeneity in ad hoc sensor networks.

Kumar *et al.* consider the suitability of two hardware platforms for various tasks, given their respective power consumption and processing ability [8]. They demonstrate that

* Other names and brands may be claimed as the property of others.

✧ The work of these authors was supported by NSF awards CNS-0325014 and CNS-0125728.

the Mica mote [2], which uses very little power but performs complex calculations slowly, is best suited to sensing, while the IPAQ, which consumes significantly more power but performs computations relatively quickly, is best suited for data fusion. They further evaluate the balance between on-sensor processing and remote processing, factoring in the cost of multihop transmission. Their results suggest that intelligent partitioning of network functions across heterogeneous nodes can increase network lifetime.

Previous analytical work has considered the amount of heterogeneity that is desirable in an ad hoc network. Mhatre *et al.* [11] evaluated a hierarchical network with two types of nodes: sensors deployed with intensity λ_0 , and nodes with higher energy and communication capacity deployed with intensity λ_1 . They prove that lifetime is maximized when λ_1 scales with the square root of λ_0 . This result is compatible with other analytical studies of hierarchical ad hoc network capacity [9].

In this paper, we utilize energy-aware routing (EAR) to help leverage energy heterogeneity. Previous work has typically focused on spreading the forwarding burden to increase network lifetime. While resulting paths tend to include low-cost links and nodes with the most remaining energy, the primary goal is load balancing and not heterogeneity. EAR+A [17] explicitly biases routes toward nodes with greater energy capacity. These nodes altruistically forward a disproportionate load. In a similar approach, a delay added to route discovery biases routes *away* from nodes with low energy capacity [1]. While we have chosen a specific solution in this paper, any energy-aware routing protocol might be used in a heterogeneous network.

Topology control protocols can also bias activity toward nodes with greater resources. GAF [18] uses geographic information to eliminate redundant nodes from sensing and communication tasks. CEC [18] has a similar goal but uses radio propagation, rather than geographic information, to identify redundant nodes. Both of these algorithms spread energy consumption over many redundant nodes. However, because node selection is biased according to available resources, it can also be used to leverage heterogeneity. ReOrg, which is similar to CEC, has been shown to be able to leverage line-powered nodes to increase network lifetime [3].

Beyond this work, key questions remain: *where, how much, and what kinds* of heterogeneity are most useful?

III. THE PROMISE OF HETEROGENEITY

In this section, we analyze the potential benefit of energy and link heterogeneity on a network. With energy heterogeneity, some nodes have no energy constraint. We call these nodes *line-powered* nodes. In link heterogeneity, some nodes have a high-quality link to the sink, and these *backhaul links* use a channel *orthogonal* to the sensor network radios. For simplicity, we consider Manhattan grid networks, in which all nodes are deployed in a grid pattern, and radio propagation is limited to immediate neighbors along vertical and horizontal axes. While not very realistic, this model is intended to provide initial insight into the potential benefit of heterogeneity.

The model also has an extremely low density, which suggests a lower bound on the benefit of heterogeneity. We back these analytical results with both experimental and simulation results in Section VI.

A. Energy Heterogeneity

In a multihop sensor network with many-to-one delivery, nodes nearer the sink expend more energy. Figure 1 illustrates this effect using data from our testbed. To increase network lifetime, therefore, nodes that consume the most energy should be line powered.

We apply the following approach to maximizing network lifetime by providing line power to a subset of nodes (the optimality of this construction is proved in Lemma 2 in Section IV.A):

- Line-powered nodes form a tree, rooted at the sink.
- The tree forms a dominating set (a dominating tree), so that every battery-powered node is one hop away from a line-powered node.

Given a sufficient number of line-powered nodes, packets can be delivered to the sink without being forwarded by any battery-powered node, increasing network lifetime. For example, in a Manhattan grid the dominating tree requires one third of the nodes to be line powered (Figure 2). Higher density networks require fewer line-powered nodes (Figures 2b and 2c).

Theorem 1: Assuming communication is the main component of lifetime, a dominating tree of line-powered nodes rooted at the sink can increase network lifetime (defined as time to first node death) by at least $\frac{n \cdot S_{e2e}}{m \cdot S_{link}}$, where n is network

size, m is the number of nodes within radio range of the sink, S_{e2e} is the average end-to-end delivery rate, and S_{link} is the link success rate in the vicinity of the sink.

Proof: Assume that each node sends packets to the sink at a




Figure 1. Energy consumption in a 5x6 testbed network with a single sink.




Figure 2. Optimal line-powered node placement in a grid topology. Grey nodes are line-powered. The sink is node S at the bottom of each network.

constant rate $1/t$ and that each radio transmission consumes cost e at each battery-powered node.

Without energy heterogeneity: One of the m nodes forwards the most packets and will limit lifetime. Since in each time period t , $\frac{n \cdot S_{e2e}}{S_{link}}$ packets are forwarded by the m nodes

one hop from the sink, a maximum network lifetime proportional to $\frac{m \cdot S_{link}}{n \cdot S_{e2e} \cdot e}$ is achieved when the load is balanced

across the m nodes. **With energy heterogeneity:** Each battery-powered node transmits one packet per time period t . All battery-powered nodes die simultaneously after a lifetime proportional to $1/e$. Thus, the lifetime improvement is at least

$$\frac{n \cdot S_{e2e}}{m \cdot S_{link}} \cdot \square$$

As an example of this result, energy heterogeneity in a 24-node network with an end-to-end delivery rate of 80%, in which 3 nodes are adjacent to the sink will multiply network lifetime by a factor of at least 6.4. Note, because packet loss (without retransmissions) only increases network lifetime in the homogeneous case (see the expressions in Theorem 1), *the apparent benefit of network heterogeneity decreases as loss rates increase.*

While energy consumption can be significantly reduced by deploying a dominating tree of line-powered nodes, some benefit can be obtained with fewer nodes. Figure 3 shows the average energy consumption of battery-powered nodes running SMAC [19] in a heterogeneous 6x8 Manhattan grid, for various percentages of line-powered nodes (deployed as in Figure 2a). Due to edge effects, the maximum benefit is achieved with slightly more than 33% line powered nodes. With 100% link success rate, more than a 70% savings in average energy is achieved. When link-layer retransmissions are enabled, energy consumption in the heterogeneous network is significantly greater, resulting in a roughly 85% savings in the heterogeneous case.

B. Link Heterogeneity

Link heterogeneity reduces the average number of hops that data packets take from each sensor to the sink. Since sensor network links tend to have low reliability, each hop significantly lowers the end-to-end delivery rate. Backhaul links

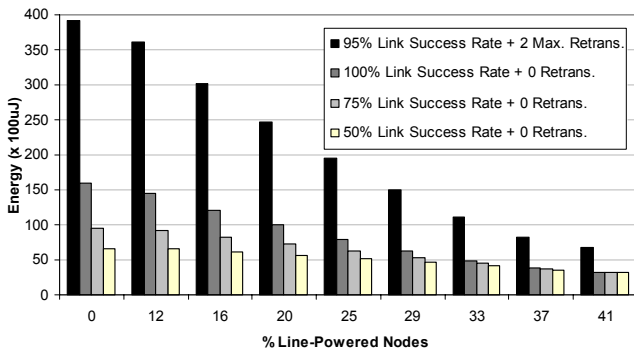


Figure 3. Analytical result of average energy consumption of a 6x8 Manhattan-grid network with various link success rates.

provide a highway bypass across the network, simultaneously increasing the end-to-end delivery rate and decreasing energy consumption.

In an $m \times n$ Manhattan grid (Figure 4), the length of the shortest path from a node at position (i, j) to a sink that is adjacent to the midpoint of one edge is given by

$$d_{i,j} = \left\lfloor \frac{m-1}{2} - i \right\rfloor + (j+1) \quad (1)$$

If we consider each backhaul to be one hop from the sink, the length of the shortest path from a node to the sink via a backhaul node at location (k, l) is given by

$$d_{i,j}^{b(k,l)} = |k - i| + |l - j| + 1 \quad (2)$$

Thus, the delivery path length for a node at (i, j) given the choice of one or more backhaul links is $d_{i,j}^*$, the minimum of $d_{i,j}$ and $d_{i,j}^{b(k,l)}$ for each available backhaul node. The resulting sum of the shortest path lengths from each node to the sink is given by

$$S_{m,n} = \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} d_{i,j}^* \quad (3)$$

For a fixed number of backhaul nodes, there exists a deployment that minimizes $S_{m,n}$. We refer to this placement of nodes as the *optimal placement* of the backhaul links for a given $m \times n$ network.

Using Equation 3, Figure 5 shows the coordinates of an optimal location of a single backhaul node, relative to network shape, $(X_0/X, Y_0/Y)$. While a specific fixed Y value was used, an identical result can be obtained in any grid network. If X and Y are the network width and depth respectively, when $X/Y = 1$ the network is square. For $X/Y < 1$, the network is deep (from the sink's perspective), and for $X/Y > 1$, the network is broad.

In deep or square networks, optimal placement of the backhaul node is always at $(X/2, 2Y/3)$. To see why, divide the network into three equal parts along the Y axis. Packets from the two thirds furthest from the sink pass through the backhaul node. Packets from the remaining third do not.

When $X/Y > 3$, the optimal location for the backhaul node is $2/3$ of the distance from the sink along the X axis with a Y coordinate that asymptotically approaches the center of the axis. Intuitively, when $X \gg Y$, it is best to place the backhaul node in the right or left half of the network. In the extreme case, the network becomes a line along the X axis, with a




Figure 4. An $m \times n$ network with a single sink (S).




Figure 5. Optimal placement of one backhaul node at (X_0, Y_0) in a grid network of depth $Y=20$ and a varying width X . This sink node lies at the midpoint of the X axis (as shown).

backhaul link $2/3$ of the distance from the sink to one end of the X axis.

To evaluate optimal deployments identified using Equation 3, Figure 6 presents average path length in grid networks of 180 nodes with varying dimensions and various numbers of backhaul links. Note that a single backhaul node is more effective when $X/Y=1/2$ (a tall, thin network) than when $X/Y=2$ (a short, fat network). This occurs because the optimal location for the backhaul node is offset and thus benefits only half of the network. This effect is significantly reduced when two backhaul nodes are deployed because each serves a different half of the network. As the number of backhaul nodes increases further, the average path length flattens dramatically over all network shapes.

Decreased path length provides a corresponding increase in end-to-end success rate. Figure 7 shows the average path length and data delivery rates in a 7×8 Manhattan-grid network with increasing numbers of optimally placed backhaul nodes. The addition of just one backhaul node reduces average path length by 48%, with a lower incremental reduction from each additional node. For a link success rate of 70%, the end-to-end success rate triples from 12% to 37% with one backhaul node and nearly quadruples to 44% with three.

In a general network, the maximum benefit occurs when each node is within one hop of a backhaul node or the sink; thus end-to-end success rate is nearly equal to the link success rate. As average path length increases, the end-to-end success rate decreases. While the incremental benefit of backhaul nodes is lower in larger networks, it increases as network density increases.

IV. OPTIMAL RESOURCE DEPLOYMENT

The previous section described the benefit of carefully placed heterogeneous resources in a grid topology. This section explores the question of optimal placement in an *arbitrary topology*. We prove that the optimal placement of line-powered nodes is NP-hard but that the placement is a tree rooted at the sink. On the other hand, we show that the optimal placement of backhaul links is easily obtained by using shortest path algorithms.




Figure 6. Average path length as a function of network shape for various numbers of backhaul nodes in a 180-node grid network.

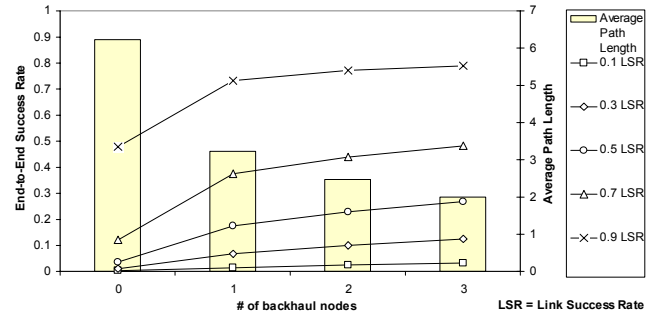


Figure 7. Analytically derived average path length and end-to-end success rate of a 7×8 Manhattan-grid network with various link success rates.

A. Energy Heterogeneity

A sensor network can be represented as a connected graph $G = (V, E)$, where the vertex set includes one designated *sink* node. Let $|V| = n + 1$, and denote the n nodes (excluding the sink) as $\{a_1, \dots, a_n\}$. Let $V' = \{b_1, \dots, b_m\} \subseteq V$ be the neighbors of the sink node. Let us initially assume that there are no line-powered nodes. The *routing* problem of interest is to find routes from all sensor nodes to the sink such that we *maximize* the time to first node failure. To quantify this metric, we attach two weights to each node that are defined as follows:

- $r(a_i)$ is proportional to the rate at which new traffic is generated by a_i . For the following discussion we assume that $r(a_i)$ is a constant.
- If T is a spanning tree rooted at the sink node then we define $w(a_i, T) = \sum_{j \in t} r(a_j)$, where t is the subtree of T rooted at node a_i .

Node lifetime is then inversely proportional to its weight w . Thus, our routing problem is the following:

Find a spanning tree T such that it minimizes $\max_{i=1}^n w(a_i, T) \triangleq W(T)$

Unfortunately, this problem is NP-hard because a solution would enable a polynomial time solution to the *maximum balanced connected partition* problem, which is known to be NP-hard. The maximum balanced connected partition problem can be stated as follows: For a connected graph $G=(V,E)$ and a non-negative vertex weight function $w:V \rightarrow \mathcal{N}$, consider a parti-

tion of the vertex set into two non-empty disjoint sets V_1 and V_2 such that the subgraphs of G induced by V_1 and V_2 are connected. The balance of the partition is defined as $B(V_1, V_2) = \min\{w(V_1), w(V_2)\}$, where $w(V_k) = \sum_{v \in V_k} w(v)$.

Thus we need to find a partition that *maximizes* B .

Theorem 2: Minimizing $W(T)$ is NP-hard.

Proof: We begin with the following observations:

- The weight of each neighbor $w(b_i, T)$ of the sink is equal to the sum of node weights $r(a_j)$ in the subtree rooted in b_i for the tree T .
- In any spanning tree T rooted in the sink, the node with greatest weight is a neighbor of the sink.
- If we assume that the sink node has exactly two neighbors b_1 and b_2 , then minimizing $W(T)$ implies minimizing $\max\{w(b_1, T), w(b_2, T)\}$ for all trees T . Equivalently, this corresponds to *maximizing* $\min\{w(b_1, T), w(b_2, T)\}$ for all trees T .

If we had a polynomial time solution to our problem we could use it to find a polynomial time solution for the maximum balanced connected partition problem as follows. Given a graph $G = (V, E)$ add a new node called the ‘‘sink.’’ Connect the sink to two nodes b_1, b_2 in V and find the spanning tree T that minimizes $B(b_1, b_2) = \max\{w(b_1, T), w(b_2, T)\}$. Repeat this for all $n(n-1)/2$ selections of the nodes b_1, b_2 . Then take the particular selection that minimizes B . The tree T produced for this partition is actually two vertex-disjoint subtrees, one rooted in b_1 and one in b_2 . The vertices in each of these subtrees form the two optimal vertex sets V_1 and V_2 that achieve the maximum balance. \square

As discussed in Section III.A, some nodes can be made line-powered (LP). LP nodes do not require a battery and will never die¹. If b is an LP node, then $w(b, T) = 0$ for all trees T . However, if node p lies on the path from b to the sink in tree T , then $w(p, T)$ is unchanged whether b is LP or not. Some interesting questions that arise are the following:

1. If we can make k nodes LP, can we identify nodes a_i that can be made LP and a corresponding spanning tree T , such that we maximize time to first node failure? In other words, minimize $W(T, k)$.
2. For a given integer $c > 0$, can we minimize the number of LP nodes k and identify the corresponding optimal tree T , such that $W(T, k) = c$? Note that $c = 0$ implies that all nodes must be LP.

We explore the above two questions below.

Lemma 1: The optimal location for a single LP node is at node b_i , i.e., an immediate neighbor of the sink.

Proof: To prove this, we observe the following. For any tree T , there is some node $b_q \in V^*$ such that,

$$W(T) = w(b_q, T)$$

In other words, the first node to run out of power is always a neighbor of the sink. This is easy to see because, for all

$a_p \notin V^*$ there is a node b_l in tree T that lies on the path between a_p and the sink. Thus,

$$w(b_l, T) > w(a_p, T)$$

Since the above observation is true for all trees, minimizing $W(T)$ by making a node LP is only achieved by making some neighbor of the sink LP. \square

Lemma 2: If we can make k nodes LP, then an optimal placement of these nodes forms a subtree rooted at the sink.

Proof: Let T be a spanning tree rooted in the sink. Let P be a placement of k LP nodes. Let $W(T, P)$ define the weight of the tree with placement P . Assume there is an edge $a_p - a_q$ in T such that a_p is LP while a_q is not, and a_q lies on the path from a_p to the sink. Finally, assume a_r is the non-LP node in tree T with placement P with the greatest weight.

Construct a new placement P' in which we make a_r LP instead of a_p . We contend that,

$$W(T, P) > W(T, P')$$

This can only be false if $w(a_p, T, P') > w(a_r, T, P)$. Clearly this inequality is not true because

$$w(a_p, T, P') < w(a_q, T, P') = w(a_q, T, P) \leq w(a_r, T, P)$$

This proves the lemma. \square

Theorem 3: Finding an optimal placement of k LP nodes that minimizes $W(T, k)$ is NP-hard.

Proof: If we set $k = 0$, we reduce the problem of finding $W(T, k)$ to finding T such that $W(T)$ is minimized (which is NP-hard by Theorem 2). \square

Theorem 4: For $c > 0$, and a given $0 < k < n$, finding a placement for k WP nodes and a corresponding tree T such that $W(T, k) = c$ is NP-hard.

Proof: If $c = 1$, the problem is identical to the problem of finding the *maximum leaf spanning tree* where the goal is to maximize the number of leaves. In such a tree, we make all nodes except the leaves LP and thus achieve our objective of $W(T, k) = 1$. Alternatively, we can also map our problem to finding the *minimum connected dominating set*. After we find this set, we find a spanning tree of the nodes that belong to the dominating set and make them all LP. Nodes outside this set are thus one hop away from the LP nodes. For $k > 1$, we can add links between nodes $\leq k$ hops apart in the graph. Solving the problem is now equivalent to the minimum connected dominating set problem. \square

B. Link Heterogeneity

A sensor network can be represented as a connected graph $G = (V, E)$. We assign each edge $e = (u, v)$ in E a weight $w(e)$ according to the reliability of packet transmission from u to v . For any spanning tree T rooted at a_i we define

$$w(a_i, T) = \sum_{e \in T} w(e)$$

Thus, our routing problem can be stated as follows:

Find a spanning tree T such that it minimizes $w(a_i, T)$

If the success rate of link l is defined to be P_l , we can assign weights such that

$$w(T) = \frac{\ln(P_l)}{\ln(\Delta)}$$

¹ For routing purposes, converting one node into an LP has the net effect that all links leaving this node are now considered to have a link weight of zero.

where Δ is the minimum possible edge weight. Dijkstra's algorithm can be used to find a spanning tree T rooted at the sink such that $w(a_i, T)$ is minimized and the average packet delivery rate to a_i is maximized.

As stated in Section III.B, a subset V' of nodes in V can be given highly reliable links to s . Thus, for each v_i in V' , there exists an edge $e_j = (v_i, s)$ such that $w(e_j) = \Delta$. Such links are deemed *backbone* links.

Theorem 5: An $O(|V|^{2+k})$ algorithm exists that selects k backbone links to G such that $w(s, T')$ is minimized.

Proof: For a given set of k nodes, k edges can be added to G and assigned weight Δ . Dijkstra's algorithm can be used on the resulting graph to find the spanning tree T with minimal $w(s, T)$ in $O(|V|^2)$ time. Since there are less than $|V|^k$ selections of k nodes from V , all possible combinations can be checked in $O(|V|^{2+k})$ time. \square

V. LEVERAGING HETEROGENEITY

Exploiting heterogeneity requires careful placement of these resources (as described in Sections III and IV) and the design of protocols that can utilize these resources for maximum benefit. In this section we examine two key protocol elements: resource-aware routing and heterogeneous MAC design.

A. Energy Heterogeneity

The following subsections describe mechanisms to exploit line-powered nodes in an ad hoc sensor network.

1) Path Selection and Energy Heterogeneity

The goal of path selection in a network with energy heterogeneity is to encourage packet flow to and along branches of a line-powered dominating tree, where the cost of forwarding is zero. In the absence of a dominating tree, routing should be biased to select paths that leverage line-powered nodes. Any energy-aware routing protocol can find these paths, given an appropriate metric (as in Equation 4). We use single-destination DSDV [12] because code for the mote platform was available.

The energy metric used for route selection captures the end-to-end energy cost of delivering a packet to the sink. The link cost is based on the expected energy cost to transmit, receive, and retransmit a packet. For nodes that are line-powered, transmission and reception of packets has zero cost, which is reflected in the metric.

Each DSDV route update (*RUPDATE*) contains a metric estimating the sender's cost to deliver a packet to the sink, as well as the sending node's energy source (battery-powered or line-powered). At each hop, the *RUPDATE* receiver adds to the received metric the cost associated with its link to the sender. The resulting metric is used to select the optimal next hop to the sink.

To compute link cost, we assume a CSMA/CA protocol that utilizes a RTS/CTS/DATA/ACK exchange. Energy cost for reliable transmission of a data packet over a single hop can be computed as follows:

$$E_{link} = \max \{ ((E_{tx-rts} + E_{rx-cts} + E_{tx-data} + E_{rx-ack}) * b_t + (E_{rx-rts} + E_{tx-cts} + E_{rx-data} + E_{tx-ack}) * b_r) * R_l, \Delta \} \quad (4)$$

Above, E_{tx} and E_{rx} are transmission and reception costs respectively. The variables b_t and b_r are set to 1 if the transmitter or receiver (respectively) are battery-powered, 0 otherwise. R_l is the expected number of retransmissions, which is computed as follows:

$$R_l = \frac{1 - (1 - P_s)^N}{P_s} \quad (5)$$

Above, P_s is the link success rate, and N is the maximum number of retransmissions before the packet is discarded. If both the sender and receiver are line-powered then the first term in Equation 4 is zero. In order to avoid a link energy cost of zero we compute the maximum of the true link cost and a small constant Δ .

Note that this metric will always select paths that include as few battery-powered receptions and transmissions as possible, and will therefore utilize a line-powered dominating tree, if one exists.

While this metric minimizes total energy consumption it does not necessarily maximize network lifetime or enable networks of nodes with heterogeneous battery capacities. We can address this concern by replacing raw energy with remaining node lifetime in our metric. This metric can balance the load across available paths over time, accounting for energy drain of individual nodes. Also, this metric can exploit individual nodes with greater energy capacity, in addition to line-powered nodes with infinite energy capacity, until that extra capacity is consumed. In Section VI, we have chosen to evaluate the simpler metric based on energy consumption, which will provide equivalent results in the relatively short period of our experiments.

2) S-MAC and Energy Heterogeneity

The MAC layer plays a key role in extending the lifetime of a battery-powered sensor network. Ye *et al.* identify four major sources of communication-related energy consumption: collisions, overhearing, idle listening, and control overhead [19]. An energy-conserving MAC avoids these wasteful activities, allowing nodes to sleep as much as possible. A MAC layer designed to exploit heterogeneity can *further decrease the duty cycle of battery-powered nodes* by shifting the burden to nodes with greater energy capacity.

We use S-MAC [19] as the basis for our heterogeneous MAC protocol, because of its ready availability for simulation and implementation. S-MAC uses periodic listen and sleep schedules, and neighbors exchange schedules, forming a virtual cluster of nodes that wake and sleep together. S-MAC also includes an *adaptive listening* mechanism, to reduce latency caused by periodic sleeping in a multi-hop network. Nodes that overhear an RTS/CTS exchange wake up for a short period at the end of the transmission, allowing the receiver to immediately forward the packet along a multihop path, rather than waiting for its next scheduled wakeup time.

We modified S-MAC to exploit energy heterogeneity, allowing battery-powered nodes to save additional energy. Our modifications to S-MAC are outlined below:

1. Because nodes that do not forward packets do not benefit from adaptive listening, we selectively disable adaptive listening on nodes that are leaves of the routing tree. However, all nodes still use adaptive listening to retransmit unacknowledged transmissions, which reduces latency.
2. Nodes that are leaves of the forwarding tree do not receive data packets. Since route updates and S-MAC synchronization packets are broadcast packets, they are received at the start of the listen period. Thus, leaf nodes are assigned a shorter listening period, further reducing idle listening.
3. Synchronization packets are sent more frequently by forwarding nodes than by leaf nodes. Thus, a greater synchronization burden is placed on forwarding nodes, which tend to be line powered.

Similar modifications could be made to other energy-conserving MACs, such as TRAMA [13] or T-MAC [15], allowing them to leverage heterogeneity as well. For example, schedule-based MACs, like TRAMA, could place more of the scheduling burden on nodes with greater resources.

B. Link Heterogeneity

The following subsections describe mechanisms to exploit reliable backhaul links in an ad hoc sensor network. As in the previous section, we consider issues of path selection and medium access control.

1) Path Selection and Link Heterogeneity

Backhaul nodes connect the sensor network to another network at two or more locations, forming a tiered architecture (Figure 8). The upper-layer network forms an overlay that allows high-bandwidth, highly-reliable communication between backhaul nodes and the sink. The overlay network may be an ad hoc network, or it may be part of an existing wired or wireless infrastructure. As such, the overlay may already have its own routing mechanism, which we cannot arbitrarily replace.

The goal of introducing heterogeneous links into the network is to increase the rate of successful packet delivery to the sink. If backhaul nodes are well deployed in the network (as described in Section III.B), then nodes near a backhaul node will have a choice between sending a packet over multiple sensor network hops to the sink or a few hops to the backhaul node followed by a highly-reliable hop to the sink. Metric-based routing, which has been shown to improve end-to-end




Figure 8. A hierarchical network architecture and a subset of packet paths.

path quality in homogeneous networks [4], can also provide the basis for determining whether a given node should utilize the overlay network to deliver a packet to the sink.

We use a single-destination form of DSDV with an end-to-end reliability metric for routing (as described in [20]) which is similar to the One Phased Pull mode of Directed Diffusion [5]. Using a link-layer sequence number, each node maintains a history of the reception success rate for each neighbor (sequence number gaps indicate a missed packet). Nodes share this information in a periodic broadcast and calculate link quality as the minimum success rate in each direction. Link quality is combined end to end to select paths with the greatest likelihood of delivery to the sink.

To integrate the sensor network with the overlay network, we developed a transparency layer (Figure 9) which hides the presence of multiple interfaces from the routing layer. The transparency layer learns, by listening, which nodes are reachable on each interface, and forwards each message on the correct interface. The node’s radio is used to send packets to true neighbors, and the overlay interface is used for neighbors adjacent via the overlay network. Broadcast packets are sent on both interfaces. Thus, incoming packets from either interface are indistinguishable by the routing layer.

In our testbed, backhaul nodes are actually two nodes coupled by a serial link: a sensor node and an embedded PC attached to an IP-based network. A process running on the embedded PC encapsulates packets received over the serial link into IP datagrams and sends them to the sink. Packets received over the IP network are decapsulated and delivered to the sensor node via the serial link.

Routing in the overlay and sensor network are completely independent (Figure 8). Encapsulated packets are routed on the overlay as IP payload. Route update packets leaving the sink node are forwarded both by the sensor network radio and across the overlay to all “adjacent” backhaul nodes. Because the overlay is highly reliable, the backhaul nodes measure high reliability on this “link” and advertise a favorable metric to their neighbors. Upon receipt of a route update, the transparency layer on a backhaul node notes that the sink is a neighbor over the backhaul interface. Thus, data packets forwarded by the backhaul node will utilize the overlay. No routing layer integration is required.

2) S-MAC and Link Heterogeneity

In exploiting heterogeneous links, our goal is to increase the end-to-end delivery rate of the network. One limitation of sensor nodes is that they tend to have limited memory capac-




Figure 9. Hierarchical networking software architecture.

ity. While our backhaul nodes can have virtually unlimited packet buffering, most sensors have very small packet buffers. So, we modified S-MAC to make it more resilient to buffer overflows.

If a node's in-bound buffer is full, it can either accept the data packet and eventually drop it, or inform the sender to refrain from transmission. We opt for the second alternative; a node receiving an RTS will reply with an *N-CTS* if its in-bound buffer is full, indicating that a packet should not be sent. While not specifically related to link heterogeneity, we observe that this simple feature significantly improved end-to-end reliability.

VI. EXPERIMENTAL EVALUATION

In this section, we evaluate and verify through testbed and simulation experiments the impact of the degree and placement of heterogeneous resources in networks of various size and density.

A. Experimental Methodology

We conducted experiments using both OPNET* simulations and real testbeds. Our testbeds consisted of Mica motes (Atmel® ATmega103L AVR® microcontroller with 128 KB of flash, 4 KB of RAM, and an RFM TR1000 916 MHz radio transceiver) with a 20 Kbps (encoded) wireless channel [2]. Protocols and applications were implemented using TinyOS [6], an event-driven operating system designed to fit within the minimal resources of mote hardware. We considered $m \times n$ grid topologies of various sizes (Figure 4), where $m < n$. While a real sensor network would not typically be deployed in such a uniform fashion, a grid topology was chosen to enable application of and comparison with analytical results, as well as to allow a repeatable performance evaluation.

Most experiments were conducted at approximately unit density. Transmit power was set such that nodes could communicate only with their immediate neighbors. In experiments in which density varied, the physical dimensions of the network were halved for a density of 1:2 and divided by three for a density of 1:3, while the number of nodes was held constant.

Due to the challenge of measuring extremely small energy consumption over many nodes, energy consumption was estimated using a mathematical model based on time actually spent in various operating modes. Energy heterogeneity was realized in software, simply by not considering the energy consumption of the nodes designated as line-powered. Link-heterogeneity was realized by attaching each selected node to an embedded PC, which provides access to the IP-based backhaul network. The backhaul network is implemented using an Ethernet LAN. This network approximates an 802.11 network, in that it has a packet delivery rate and bandwidth orders of magnitude greater than the mote network. We believe that this approximation will not affect our results. Except where noted, line-powered and backhaul nodes were placed in the grid according to the strategy described in Section III.

Most of our experiments utilized optimal deployments of line-powered or backhaul nodes in grid networks of various densities. To achieve a minimum dominating tree (as discussed in Section III.A) line-powered nodes were deployed as

shown in Figure 2. Optimal deployments of backhaul nodes in a grid network (Table I) were determined as described in Section III.B and are independent of density.

In all experiments, all nodes initiated one packet at a regular time interval. Timing was randomly offset across nodes to avoid synchronization. All packets were destined for a single sink node located at the midpoint of the narrow edge of the network (Figure 4). Each packet was 37 bytes, including headers and payload. A low data rate was used to avoid congestion.

Simulation studies are also presented to complement and extend the testbed results, allowing factors such as latency to be evaluated. The protocols described in Section V, including our modified S-MAC and DSDV protocols, were implemented in OPNET*. Our node model is based on the Mica mote, with limited buffer capacity of one in-bound and one out-bound packet. We use a free space propagation model with a transmit range as described above and a channel error rate that achieves a 95% packet success rate in energy heterogeneity experiments (Section VI.C) and 85% in link heterogeneity experiments (Section VI.D). A maximum of three retransmissions were allowed per packet transmission. The carrier sense threshold is set to twice the reception range.

Multiple trials were used in all experiments, and all results indicate 90% confidence intervals. The duration of each experiment trial was at least 1200 seconds.

B. Evaluation Metrics

Energy heterogeneity decreases the energy that nodes consume forwarding packets and thereby increases network lifetime. We evaluate the effectiveness of energy heterogeneity using two metrics:

- Percentage reduction in average per-node energy consumption of battery-powered nodes, normalized over goodput.
- Percentage increase in network lifetime normalized over goodput.

We define lifetime to be the time to first node death. We normalize each metric above to network goodput, to avoid giving unfair advantage to less reliable networks.

Link heterogeneity increases the rate at which packets are successfully delivered end to end by decreasing the depth of the routing tree. Thus, the length of individual delivery paths

TABLE I
OPTIMAL LOCATIONS (ROW, COLUMN) OF BACKHAUL NODES FOR LINK HETEROGENEITY EXPERIMENTS.

		Network Dimensions		
		7x8	5x6	3x4
Number of Backhaul Nodes	1	(3,5)	(2,4)	(1,2)
	2	(1, 5), (5, 5)	(1, 4), (4, 3)	(0, 1), (1, 3)
	3	(1, 3), (3, 6), (5, 3)	(1, 2), (1, 5), (4, 3)	(0, 1), (1, 3), (2, 1)
	4	(1, 3), (1, 6), (5, 3), (5, 6)	(0, 4), (1,2), (3, 4), (4,2)	(0, 1), (1, 3), (2, 1), (2, 2)
	5		(0,4), (1,2), (2,5), (3,3), (4,1)	
	6		(0,0), (0,4), (1,2), (2,5), (3,3), (4,1)	

decreases. Besides increasing delivery success rates, shorter paths also decrease latency. We evaluate link heterogeneity using three metrics:

- End-to-end success rate
- Average path length to the sink
- End-to-end latency (simulation only)

The end-to-end success rate was measured by dividing the total number of data packets successfully received at the sink by the total number of expected packet transmissions for a given time period and number of nodes (since nodes transmit at regular intervals). In simulation only, latency was measured from the moment the packet arrived at the MAC layer of the originating node to the moment it was received by the application layer on the sink.

C. Energy Heterogeneity

In this section we evaluate the benefits of degree and placement of line-powered (LP) nodes in networks of various sizes and densities.

1) Effect of Number of Heterogeneous Nodes

Section III.A predicts an incremental benefit to average per-node energy consumption and network lifetime as the number of optimally-placed LP nodes increases, up to a maximum deployment of 33% LP nodes. We measure this effect in a 6x8 network using simulation and testbed experiments. Figure 10 and Figure 11 show the energy savings observed in simulation and testbed, respectively. Energy savings in simulation and the testbed differ because the end-to-end delivery rate is lower in the testbed. Figure 3 shows the correlation between packet delivery rate and expected energy savings.

In general we observe continuous improvement of energy saving as we increase the number of the LP nodes. With 33% of LP nodes a maximum energy savings of 83% was achieved in simulation. This result correlates with the expected value from Figure 3 with a 95% link success rate and a maximum of three retransmissions, as specified in the simulation. In addition, the incremental benefit decreases as more LP nodes are added, suggesting that a relatively large benefit can be achieved with a partial deployment.

The testbed results differ significantly from the simulation results. With 33% of LP nodes we achieved an energy savings of 50%, which correlates with the expected value for a 75% link success rate, as measured in the testbed. However, with fewer LP nodes, significantly less benefit is achieved. We attribute these results to the lower link success rate and lack of link-layer retransmission in the testbed. During our experiments, a significant number of non-LP nodes used non-optimal routes, most likely due to lost routing packets. This effect is much greater when multiple hops are required to reach an LP node, as is the case with fewer LP nodes. As in the simulation experiments, the energy savings does not increase beyond 33% LP nodes (not shown).

Figure 11 and Figure 12 show the lifetime improvement observed in the testbed and simulation, respectively. Lifetime increased incrementally as we increased the number of the LP nodes. Both experiments achieved a 450% increase in network lifetime with 33% LP nodes, approximately correlating

with the result predicted by Theorem 1. Beyond 33% LP nodes, no additional lifetime gain was observed, as expected.

2) Optimal vs. Random Placement

Ideally, we would like the benefit of energy heterogeneity to be insensitive to node placement. Unfortunately, this is not the case. Randomly selecting various numbers of LP nodes produced a wide range of results (Figure 10). While some random deployments produced a significant average energy savings, none did as well as optimal deployment. In most cases, fewer optimally placed LP nodes outperformed a much larger number of randomly placed LP nodes. Random deployment is suboptimal for two reasons. First, because nodes close to the sink consume the most energy, selecting any other

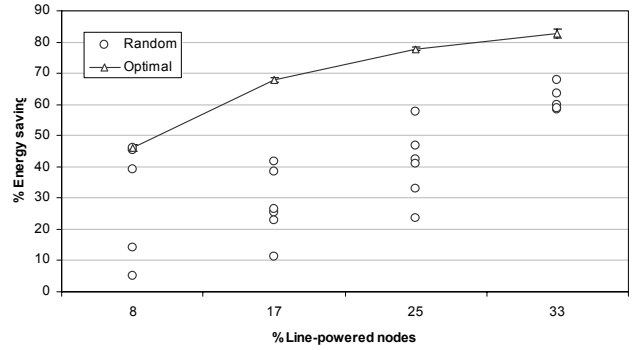


Figure 10. *Simulation result:* The effect of various % of line-powered nodes on average energy consumption in a 6x8 network for optimal and random deployments.

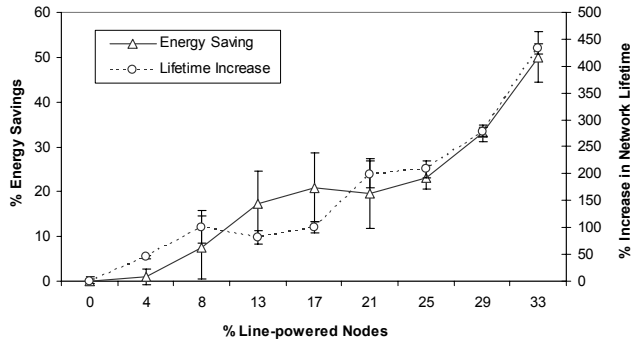


Figure 11. *Testbed result:* The effect of various % of line-powered nodes on average energy consumption and lifetime in a 6x8 network.

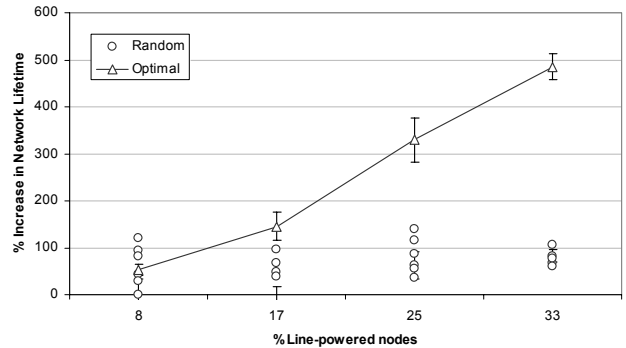


Figure 12. *Simulation result:* The effect of various % of line-powered nodes on lifetime in a 6x8 network, for optimal and random deployments.

nodes will be less beneficial. Second, by attracting packets, LP nodes increase path lengths, which increases the amount of work required at battery-powered nodes.

While random placement of LP nodes is suboptimal, the range of results appears to decrease as the number of LP nodes increases. This suggests that less care is required in LP deployment as the number of LP nodes increases.

As with average energy consumption, random deployment also fails to achieve optimal network lifetime (Figure 12). In addition, there was little improvement in lifetime as the number of randomly deployed LP nodes increased. This result supports Lemma 2. Unless there is a connected backbone to the sink, lifetime will not be significantly affected.

3) The Effect of Network Size

Figure 13 presents lifetime improvement from simulation in networks of various sizes (but constant density) with 33% LP nodes. The network lifetime increases as the network size increases. This result is predicted by Theorem 1. The slope of this line is less than would be expected purely from an increase in network size, because the end-to-end success rate also decreased as network size increased. As noted in the previous section, random deployments are not effective at increasing network lifetime. We present these results from simulation, as this allowed us greater variation in network size. We omit our testbed results, which support the same conclusions.

4) The Effect of Network Density

Figure 14 shows the effect of density on network lifetime and average energy consumption in a 6x8 grid topology. Measurements were taken at three distinct densities (1:1, 1:2, 1:3). As shown in Figure 2, at higher densities there are fewer LP nodes.

At higher density, more nodes are one hop from the sink. Thus, the lifetime of the homogeneous case is increased because traffic to the sink can be carried by a larger number of nodes. As a result, the percentage improvement of the heterogeneous case over the homogeneous case decreases as density increases (Figure 14). In addition, at higher density, the same absolute lifetime can be achieved with fewer LP nodes (33% at 1:1, 17% at 1:2, and 11% at 1:3).

Average energy consumption is also affected by network density. At higher density, the average path length is reduced. Since nodes spend less energy forwarding packets in the ho-

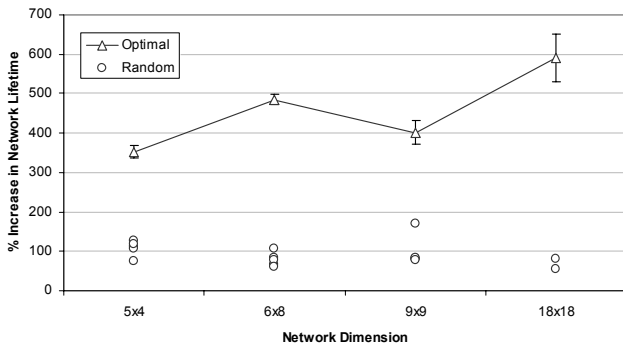


Figure 13. *Simulation result:* The effect of network size on lifetime. Both optimal and random deployments are shown.

mogeneous case, the percentage improvement decreases as density increases.

5) Summary of Results

The simulation and testbed evaluation presented in this section suggest the following conclusions:

- Optimal placement of up to 33% LP nodes can provide significant but decreasing incremental gains in network lifetime.
- Packet loss reduces the potential benefit from LP nodes, particularly with fewer LP nodes.
- Random LP node placement has a small impact on average energy consumption but almost no impact on network lifetime.
- Larger networks have a greater potential benefit from LP nodes.
- While higher density reduces packet forwarding (and thus energy consumption), fewer LP nodes are required for maximum benefit.

D. Link Heterogeneity

In this section we evaluate the benefits of degree and placement of backhaul (BH) nodes in networks of various sizes and densities.

1) Effect of Number of Heterogeneous Nodes

Section III.B predicts an incremental benefit to average end-to-end success rate as the number of optimally-placed BH nodes increases. This benefit is primarily due to a decrease in average path length. To quantify the benefit of varying numbers of BH nodes, we simulated a 7x8 grid topology and ran testbed experiments on a 5x6 grid topology.

In both simulations and testbed experiments, the average path length (end-to-end delay in simulation) decreased (Figure 15 and Figure 17) and end-to-end data success rate increased (Figure 16 and Figure 18) as the number of BH nodes increased. The addition of one BH node to a homogenous network had the greatest impact. In our simulations, we achieved a 16% decrease in the end-to-end delay and a 51% improvement in success rate. In the testbed, we achieved a 36% reduction in the average path length and a 24% improvement in the success rate. Note that these results differ slightly due to the different network sizes and link success rates used in simulation and testbed experiments.

Each additional BH node provides a gradually declining in-

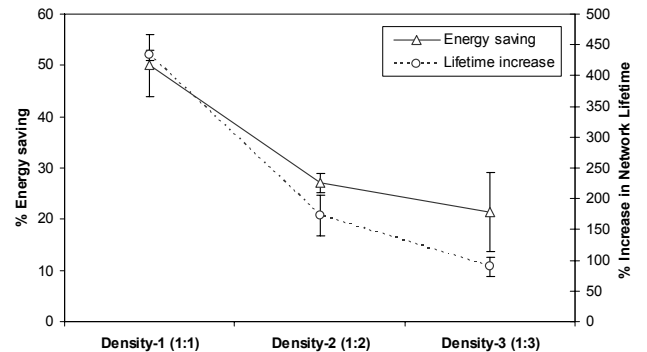


Figure 14. *Testbed result:* The effect of density on lifetime in a 6x8 network.

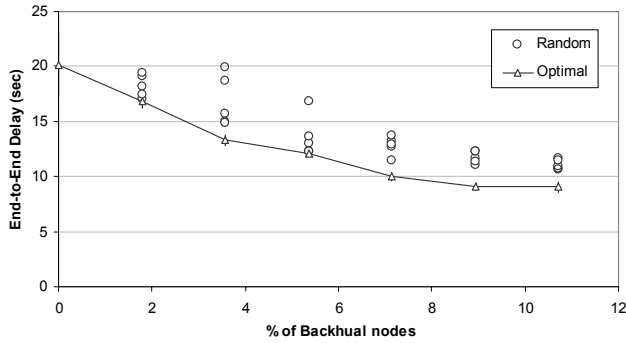


Figure 15. *Simulation result*: The effect of varying the number of backhaul nodes on average end-to-end delay in a 7x8 network. Results for optimal and random placement are shown.

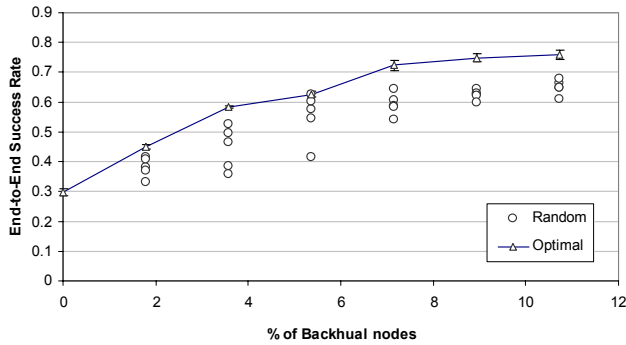


Figure 16. *Simulation result*: The effect of varying the number of backhaul nodes on end-to-end success rate in a 7x8 network. Results for optimal and random placement are shown.

cremental benefit. In simulation, we achieved a maximum end-to-end success rate of 76% (a 154% improvement) with a 55% reduction in end-to-end delay when 11% of the nodes were BH nodes. In our testbed, we achieved a maximum end-to-end success rate of 79% (a 71% improvement) with a 57% reduction in average path length with 13% BH nodes. Note that in Figure 17 and Figure 18, the maximum benefit occurs when the average path length is 2, which corresponds to all non-BH nodes being one-hop from a BH node. These results confirm that by having only a modest percentage of BH nodes, we can achieve a significant benefit, with an end-to-end success rate that asymptotically approaches the link success rate. Thus, link heterogeneity will have the most benefit to networks with deep routing trees or low link success rates.

2) Optimal vs. Random Placement

Ideally, we would like the benefit of link heterogeneity to be insensitive to node placement. We measured the performance of various numbers of randomly deployed BH nodes and compared the results with the corresponding results for optimal deployment (Figures 15 through 18).

In both simulation and testbed experiments, random deployment does not reduce path length or increase success rates as much as optimal deployment. In many cases, fewer optimally placed BH nodes can outperform a given number of randomly placed nodes. For randomly placed BH nodes, the variance in the success rate decreases as the number of BH nodes increases. Thus, careful deployment of BH nodes be-

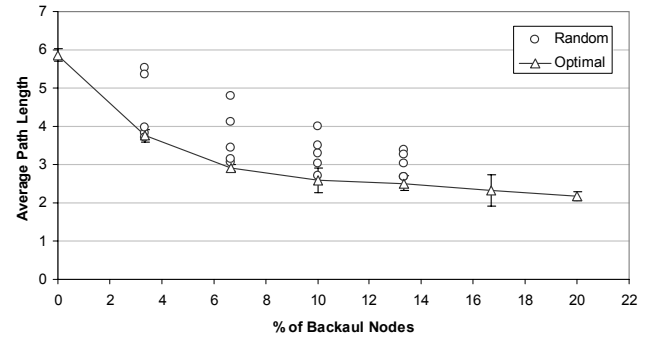


Figure 17. *Testbed result*: The effect of a varying amount of link heterogeneity on average path length in a 5x6 network. Results for optimal and random placement are shown.

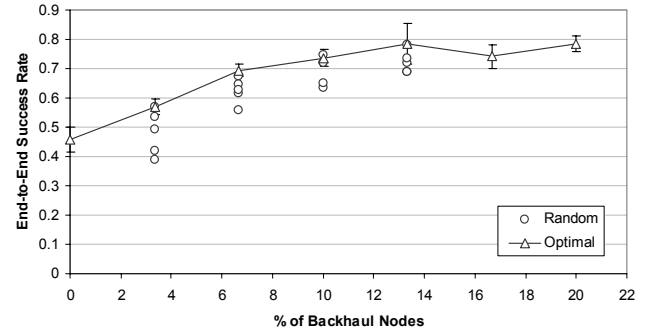


Figure 18. *Testbed result*: The effect of a varying amount of link heterogeneity on end-to-end success rate in a 5x6 network. Results for optimal and random placement are shown.

comes less important as the number of BH nodes deployed increases.

3) The Effect of Network Size

To determine the effect of network size on the benefit of network heterogeneity, we evaluated networks of different sizes (but constant density) using simulation and testbed experiments (simulation data is omitted since it is similar to testbed data). Average path length increases as network size increases, resulting in a corresponding decrease in delivery success rate. For a given number of backhaul nodes, the increase in end-to-end success rate is greater in larger networks (Figure 19). Finally, the incremental benefit of backhaul nodes diminishes as nodes are added, suggesting that no more than 10% of nodes need backhaul links, at this density.

4) The Effect of Network Density

To determine the effect of density of the network on the benefits achieved, we performed testbed experiments on a 7x8 grid with three different densities (1:1, 1:2, 1:3).

In Figure 20, for a given number of backhaul nodes, denser networks have better end-to-end success rate. It is true that as we increase the network density, collisions and channel contention increases, but the impact is smaller than the impact of decreased path length (Figure 21). In addition, at higher density, path length reaches its minimum with fewer BH nodes than at unit density. Thus, a small number of BH nodes can have a significant impact at high density.

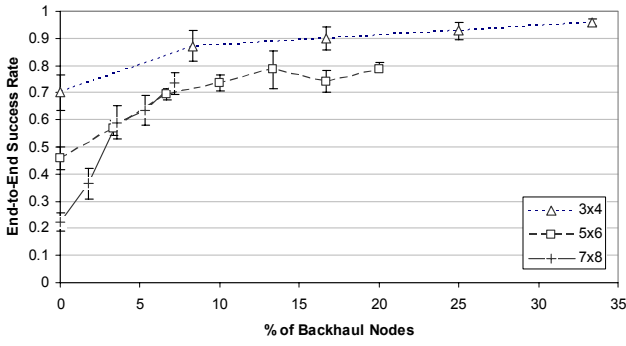


Figure 19. Testbed result: The effect of network size on end-to-end success rate.

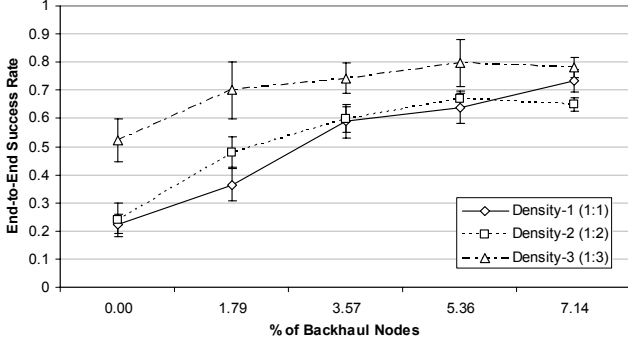


Figure 20. Testbed result: The effect of network density on end-to-end success rate in a 7x8 network.

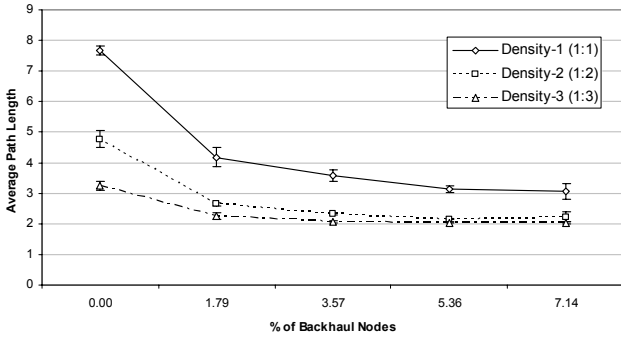


Figure 21. Testbed result: The effect of network density on path length in a 7x8 network.

5) Summary of Results

The simulation and testbed evaluation presented in this section suggest the following conclusions:

- Link heterogeneity of up to 10% can improve end-to-end success rate until it reaches the link success rate.
- Placement of backhaul nodes has much less impact on end-to-end success rate than placement of line-powered nodes has network lifetime.
- Link heterogeneity has the greatest impact in networks that are large, have deep routing trees, or have low link success rates.

VII. APPLICABILITY OF RESULTS

The results in the preceding section were obtained using specific routing and MAC protocols: DSDV and S-MAC. However, it is possible to estimate the degree of benefit that

heterogeneity can provide when other routing and MAC protocols are used, assuming a similar traffic mode.

A. Energy Heterogeneity

Energy heterogeneity increases network lifetime by decreasing the amount of work (in the form of packets forwarded), performance will be determined by the relative cost of packet forwarding and static protocol overheads. To impact of different protocol costs, we define the following parameters:

- E_{data} : the total energy spent by all nodes one hop from the sink to send and receive data packets, without heterogeneity.
- $E_{protocol}$: the total energy spent by all nodes one hop from the sink associated with per-packet protocol overhead, without heterogeneity.
- E_{idle} : the total energy spent by all nodes one hop from the sink to run an idle protocol stack (over a time period equivalent to that of $E_{protocol}$ and E_{idle}), without heterogeneity.
- E'_{data} , $E'_{protocol}$, E'_{idle} : Equivalent to the above, except energy consumption is measured per node and with heterogeneity.

These parameters can be used to determine the degree to which heterogeneity can increase the lifetime of a network, as follows:

$$\% \text{ improvement} = \frac{E_{idle} + E_{protocol} + E_{data}}{m \cdot (E'_{idle} + E'_{protocol} + E'_{data})} \quad (6)$$

In the above, m is the number of nodes that are one hop from the sink (applied as in Theorem 1) and $E_{idle} = m \cdot E'_{idle}$. By applying the values from TABLE II (obtained from our mathematical model of the energy consumption of SMAC and DSDV) and a value of $m=2$ (observed in our simulation experiments in Section VI.C.1) to Equation 6 yields a lifetime improvement of roughly 5 times, which corresponds to the results reported in Figure 11 and Figure 12.

Equation 6 can also be used to estimate the impact of other routing and MAC protocols on the benefits of energy heterogeneity. For an ideal MAC and routing layer, $E_{idle} + E_{protocol} = 0$ and $E'_{idle} + E'_{protocol} = 0$. In a network with 80% link success rate, energy heterogeneity could increase network lifetime by 6.8 times. For a worst-case MAC and routing layer, $E_{idle} + E_{protocol} \gg E_{data}$ and $E'_{idle} + E'_{protocol} \gg E'_{data}$. In this case, energy heterogeneity could increase network lifetime by 3.9 times.

Clearly the potential benefit of energy heterogeneity is dependent on the overhead of the MAC and routing protocols. In fact, the relative cost of protocol overhead is greater in the heterogeneous case, because energy heterogeneity reduces packet forwarding at line-powered nodes. As a result, protocol overhead at non-line powered nodes becomes more significant. Thus the need for the careful design of protocols is greater when energy heterogeneity is present.

B. Link Heterogeneity

The potential benefit of link heterogeneity is limited by link success rate. Link success rate can be affected by the MAC

TABLE II
ENERGY CONSUMPTION FROM A MATHEMATICAL MODEL OF SMAC
AND DSDV.

	100% link success rate, homogeneous	80% link success rate, homogeneous	80% link success rate, heterogeneous
	$E_{idle} = 3.6$ mJ	$E_{idle} = 3.6$ mJ	$E'_{idle} = 1.8$ mJ
	$E_{data} = 175.2$ mJ	$E_{data} = 35.6$ mJ	$E'_{data} = 2.6$ mJ
	$E_{protocol} = 87.5$ mJ	$E_{protocol} = 17.3$ mJ	$E'_{protocol} = 0.7$ mJ
% Data Cost	65.5%	61.9%	51.3%
% MAC + Routing Cost	34.5%	38.1%	48.7%

and routing layers, which can mask individual transmission failures and increase channel contention with management packets. If Figure 7 is applied to the link success rate apparent above a particular MAC and routing layer, the impact of various numbers of backhaul nodes can be determined.

CONCLUSIONS

We have shown analytically that optimal deployment of heterogeneous resources is computationally hard in general. In addition, the maximum benefit that can be achieved from a particular type of heterogeneity depends greatly on the shape, size, and density of the network, as well as on the overheads of MAC and routing protocols. Protocol overheads, in particular become more important as networks become more heterogeneous.

Results from our testbed and simulation experiments with a specific MAC and routing layer suggest that modest amounts of link heterogeneity can more than triple end-to-end success rates. Similarly, energy heterogeneity can provide a more than 5-fold increase in network lifetime. Assuming moderate density, this benefit can be achieved with heterogeneity at less than 10% of nodes. While careful placement is required to maximize benefit when a limited amount of heterogeneity is deployed, an increase in the number of heterogeneous resources reduces the impact of resource placement. In addition, density significantly reduces the number of resources required for maximal benefit.

While our experimental results provide valuable insight into the benefits of heterogeneity, they are only a first step. Real application deployments have more complex topologies. Future research is required to systematically evaluate the role of heterogeneity in real application deployments, including a study of other forms of heterogeneity.

REFERENCES

- [1] I.D. Chakeres and E.M. Belding-Royer, "Resource Biased Path Selection in Heterogeneous Mobile Networks," Univ. of California, Santa Barbara, Comp. Sci. Dept. Tech. Report 2003-18, July 2003.
- [2] Crossbow Technology, Inc., "Wireless Sensor Networks (product data sheet)," <http://www.xbow.com/Products/>
- [3] W.S. Conner, J. Chhabra, M. Yarvis, L. Krishnamurthy, "Experimental Evaluation of Synchronization and Topology Control for In-Building Sensor Network Applications," *Proc. 2nd Intl. Workshop on Wireless Sensor Networks and Applications (WSNA '03)*, San Diego, CA, Sept. 2003.
- [4] D. De Couto, D. Aguayo, B. Chambers, and R. Morris, "Performance of Multihop Wireless Networks: Shortest Path is Not Enough," In *Proc. of the 1st Workshop on Hot Topics in Networks*, New Jersey, Oct. 2002.

- [5] J. Heidemann, F. Silva, and D. Estrin, "Matching Data Dissemination Algorithms to Application Requirements," *Proc. of ACM SenSys*, Los Angeles, CA, Nov. 2003.
- [6] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister, "System Architecture Directions for Networked Sensors," *Proc. of ASPLOS 2000*, Cambridge, MA, Nov. 2000.
- [7] Lakshman Krishnamurthy, "Sensor Networks: Promise and Reality," invited presentation, Sensors Expo and Conference, Detroit, MI, June 9, 2004.
- [8] R. Kumar, V. Tsiatsis, and M.B. Srivastava, "Computation Hierarchy for In-Network Processing," *Proc. of the 2nd Intl. Workshop on Wireless Networks and Applications (WSNA '03)*, San Diego, CA, Sept. 2003.
- [9] B. Liu, Z. Liu, and D. Towsley, "On the Capacity of Hybrid Wireless Networks," *Proc. IEEE INFOCOM*, San Francisco, CA, April 2003.
- [10] A. Mainwaring, J. Polastre, R. Szewczyk, D. Culler, and J. Anderson, "Wireless Sensor Networks for Habitat Monitoring," *Intl. Workshop on Wireless Sensor Networks and Applications (WSNA '04)*, Atlanta, GA, Sept. 2002.
- [11] V. Mhatre, C. Rosenbert, D. Kofman, R. Mazumdar, and N. Shroff, "A Minimum Cost Heterogeneous Sensor Network with a Lifetime Constraint," *IEEE Transactions on Mobile Computing*, Jan. 2004.
- [12] C.E. Perkins and P. Bhagwat, "Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers," *Proc. of the Conference on Communications Architectures, Protocols, and Applications*, Aug. 1994.
- [13] V. Rajendran, K. Obraczka, and J.J. Garcia-Luna-Aceves, "Energy-Efficient, Collision-Free Medium Access Control for Wireless Sensor Networks," *Proc. of ACM SenSys*, Los Angeles, CA, Nov. 2003.
- [14] S. Rhee, D. Seetharam, and S. Liu, "Techniques for Minimizing Power Consumption in Low Data-Rate Wireless Sensor Networks," *Proc. of IEEE Wireless Communications and Networking Conference (WCNC '04)*, Atlanta, GA, March, 2004.
- [15] T. van Dam and K. Langendoen, "An Adaptive Energy-Efficient MAC Protocol for Wireless Sensor Networks," *Proc. of SenSys 2003*, Los Angeles, CA, Nov. 2003.
- [16] H. Wang, D. Estrin, and L. Girod, "Preprocessing in a Tiered Sensor network for Habitat Monitoring," *Proc. of the IEEE Conf. on Acoustics, Speech, and Signal Processing*, Hong Kong, China, April 2003.
- [17] A. Willig, R. Shah, J. Rabaey, and A. Wolisz, "Altruists in the PicoRadio Sensor Network," *Proc. of the International Workshop on Factory Communication Systems*, Vasteras, Sweden, Aug. 2002.
- [18] Y. Xu, S. Bien, Y. Mori, J. Heidemann, and D. Estrin, "Topology Control Protocols to Conserve Energy in Wireless Ad Hoc Networks," UCLA CENS, Technical Report 6, Jan. 2003.
- [19] W. Ye, J. Heidemann, and D. Estrin, "An Energy-Efficient MAC Protocol for Wireless Sensor Networks," *Proc. of IEEE INFOCOM*, New York, NY, June 2002.
- [20] M.D. Yarvis, W.S. Conner, L. Krishnamurthy, J. Chhabra, B. Elliott, and A. Mainwaring, "Real-World Experiences with an Interactive Ad Hoc Sensor Network," *Proc. of the International Workshop on Ad Hoc Networking (IWAHN 2002)*, Vancouver, B.C., Canada, Aug. 2002.