Throughput Characteristics of a Minimum Energy Wireless Network

Özge H. Köymen,Volkan Rodoplu,Teresa H. Meng
Center for Integrated Systems
Stanford University
Stanford, CA 94305

Abstract—**In this paper, we evaluate the performance of the Minimum Energy Wireless Network, which seeks to provide global connectivity in an ad-hoc network while maintaining overall minimum energy for communications. The properties and advantages of employing multi-hop communication over a direct peer-to-peer communication system are illustrated and the performance of such a network is characterized. The blockage percentage for throughput and the overall power consumption are also quantified via simulation.**

I. Introduction

This paper evaluates the performance of the Minimum Energy Wireless Network presented in [9], in which it was shown that the proposed network achieved the lowest power consumption for any given topology. Power is minimized by effective use of multi-hopping, thus avoiding long distance transmission. This paper examines the practicality of such a network strategy and gives insight into the design of multi-hop networks in general. This paper also includes a performance evaluation of the achievable network throughput compared to other networks and a measure of the energy overhead in forming the network topology. In addition, we propose an acknowledgement scheme for neighbor identification.

The basic idea of the Minimum Energy Wireless Network is for each node to find its neighbors to relay its packets rather than to transmit them to their destinations directly. Each node is required to enclose itself, where the enclosure of a node is a region (denoted as enclosure region) around the node beyond which it is not power-efficient for the node to transmit and thus should relay (see Figure 1). By using cost distribution, an optimal path is found to each node in the network. It was proven in [9] that strong connectivity is guaranteed using this network, allowing communication between any two nodes in the network.

Since power falls with distance as $1/d^n$, where $n \geq 2$ depending on the path loss model [8], it is advantageous not to transmit directly if relaying is an option. Figure 1 illustrates one realization of the Minimum Energy Wireless Network compared with a direct peer-to-peer network from the perspective of one user of interest.

While the power advantage is obvious, this multi-hop network has several drawbacks. The Minimum Energy Wireless Network requires an enclosure period during which nodes rediscover their neighbors and re-compute the optimal paths to their peers. A network with high mobility or unreliable nodes would require the nodes to re-enclose themselves more frequently. Re-enclosure may significantly reduce the effective data rate (throughput) available for data communications because of the time required for each re-enclosure.

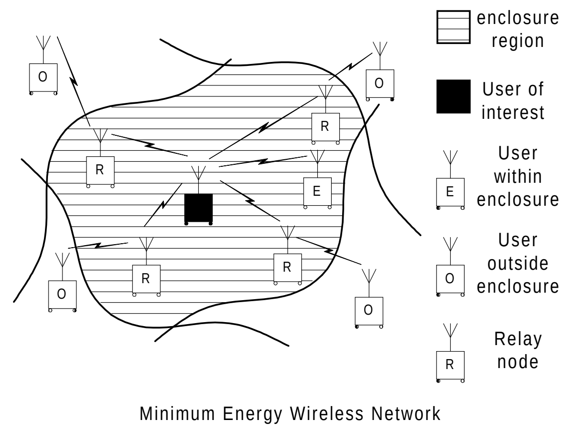

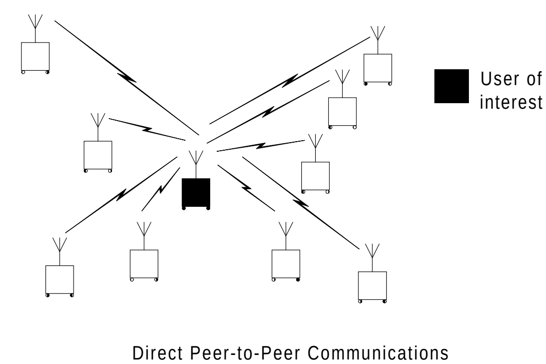

Fig. 1. Comparison of Networks

The data rate at which a node can send packets to another node over this multi-hop network is limited by the maximum data rates that can be accommodated by the relay nodes. The relay nodes in the network implicitly handle more data since they act as routers for the surrounding nodes. With a limited throughput per node, these relay nodes will present a bottleneck to the system. This assumes that the amount of bandwidth assigned to each node is the same although this problem could be alleviated with a clever bandwidth assignment scheme.

Despite these apparent limitations and drawbacks, this system provides a robust and fault-tolerant network that is ideal for applications that require absolute lowest power while satisfying a bursty data pattern and a low mobility constraint. Mobility can certainly be accommodated as long as the re-enclosure period is chosen appropriately. Due to the focus on low mobility applications, our simulations will consider only stationary nodes. We also assume orthogonal transmissions and ignore interference between users.

We begin this paper by providing a detailed overview of our system design in Section II, where we present the design assumptions and limitations. The paper continues with simulation results presented in Section III. We end the paper with concluding remarks in Section IV.

II. System Overview

A. Assumptions

The focus of this paper is the evaluation of the Minimum Energy Wireless Network under a realistic simulation environment. The following are some important assumptions that were made:

1. *Proactive* Routing - We use a proactive routing protocol; this means that a route for a packet is known a priori. The routing tables are calculated within the enclosure periods, during which neighbors are found, whose update frequency is dependent on the mobility and topology of the network.
2. CDMA - We assume that each node is assigned an orthogonal multi-access code. We do not consider interference from other nodes and assume that all nodes operate in a multi-access fashion. This allows each node to receive packets from all its neighbors simultaneously when needed.
3. Power Limitation - We assume that the nodes are not limited by transmit power. That is, a node will transmit at any power required to reach its neighbors. In a system with many nodes over the deployment region, the enclosure region for each node will be small. In this case, the transmit power limitation becomes irrelevant since nodes do not need to search for neighbors beyond a small area around themselves.
4. Cost Tables - The formulas to calculate the cost tables used are known to all nodes for optimal path selection.

B. Timing Structure

There are three distinct periods in the timing structure: *broadcast*, *cost distribution*, and *data communications*.

B.1 Broadcast Period

This period is reserved for each node to find its neighbors and enclose itself within a specified time of T_B using an acknowledgement scheme. The protocol, which is illustrated in Figure 2, is defined as follows:

1. Each node starts with an initial power and sends a broadcast message with *Transmit Power* specified in a packet field.
2. The node enters an idle state waiting for a reply.

(a) If no reply is received within a *timeout* period, another broadcast is sent with increased power. The power is increased by a pre-determined step size.

(b) If a reply is received, the node processes the packet for its contents.

3. The contents of the reply are checked and fall into two categories:

(a) Broadcast Message

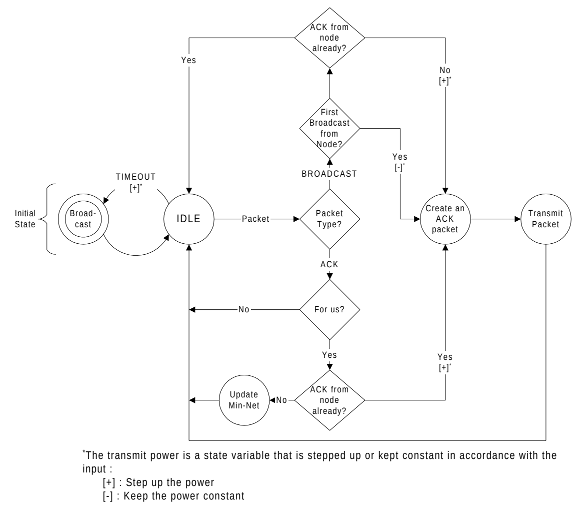

Fig. 2. Broadcast Algorithm

- If this is the first broadcast message received from a particular node, then an ACK is sent to the node with the *Transmit Power* field of the received packet set in the ACK packet as the *Received Power*.
- If the node has already received a broadcast from this node but has not received an ACK yet, then it sends another ACK with more power. Again, the *Transmit Power* field of the received packet is set as the *Received Power* field in the ACK packet.
- In all other cases, the node enters an idle state.

(b) ACK Message

- If the ACK is meant for the node and it is the first one received from the remote node, the enclosure is updated with the new node.
- If the ACK is meant for the node and it already has an ACK from the remote node, another ACK is sent with more power. The power needs to be increased since the node can only assume that its previous ACK did not have the required power to be received.
- In all other cases, the node enters an idle state.

The broadcast is acknowledgement (ACK) based; a neighbor is not considered valid until an ACK to a broadcast is received. Consequently, any two given nodes must receive a broadcast message and an ACK message before either of them considers the other as a valid neighbor. In this manner, the nodes perform a three-way handshake.

In the algorithm, the *timeout* needs to be set carefully, otherwise the nodes risk sending unnecessary ACK's to each other. The safest value for the *timeout* is the roundtrip time of the deployment region plus any processing delay, which is part of the system design. Note that a larger timeout will result in a larger broadcast period, T_B, but will avoid the above-mentioned problem. In addition the duration of the broadcast period is dependent on the power step size and the size of the deployment region.

It is important to transmit with just enough power to reach the next hop to realize the power advantage of the Minimum Energy Wireless Network. Each ACK has the *Received Power* field set to the *Transmit Power* of the broadcast message being responded to in order to inform the nodes of the power level required in reaching their neighbors. A safety margin can be accommodated by a multiplicative factor so that the power transmitted to the next hop is a percentage above the acceptable SNR.

B.2 Cost Distribution

The distribution of costs requires a much shorter period than the broadcast period. Each node initially needs to find the maximum power required to transmit to its neighbors. In other words, each node must distribute the costs with enough power to reach its furthest neighbor among the enclosing neighbors. After finding the appropriate power, each node executes the following algorithm:

1. The node creates a packet with the current cost table and broadcasts it to its neighbors.
2. The node enters an idle state waiting for a reply.
 (a) If no reply is received within a *timeout* period, another packet is sent.
 (b) If a cost packet is received, the node updates its cost table and enters the idle state again.

The cost distribution period, denoted by T_C, needs to be long enough to allow the cost tables to converge, while the *timeout* needs to be set to avoid unnecessary broadcasts. We chose the timeout and cost distribution period appropriately to allow convergence in the worst-case topology.

B.3 Data Communications

Packets are generated at each node with the inter-arrival time between packets modeled as a negative exponential with mean $1/\lambda$ seconds. The probability for generating a packet for a particular destination is uniformly distributed over all the other nodes in the simulation. Note that in general, the data communications period, denoted by T_D, would be chosen to accommodate mobility and changing network topologies. For stationary networks, the value for T_D can be picked arbitrarily.

Furthermore, each node is allocated a certain throughput, R_{max}, above which packets are dropped upon arrival. We assume that dropped packets are retransmitted or handled accordingly by an upper layer protocol.

III. Simulation

We used OPNET, an event driven simulation tool, to model and simulate the network performance. We examine blocking probabilities (percentage of packets dropped), power consumption (for both broadcast and communications periods), and enclosure time. The parameters varied are the cutoff rate (R_{max}) and the rate of packet generation ($1/\lambda$).

The deployment region was fixed at $1000m \times 1000m$ over which the nodes are distributed deterministically under three topologies: uniformly deployed nodes, medium-sized clusters (three clusters of five nodes each), and small clusters (five clusters of three nodes each). The nodes are allowed to transmit simultaneously free of interference from other nodes in accordance with the orthogonal CDMA assumption. It is assumed that a bit error rate of 10% is correctable by forward error correction codes. Each user is also assumed to use BPSK modulation through which the necessary SNR can be found. Each node transmits at a power level in accordance with this SNR.

The path loss is assumed to have a distance dependence of $1/d^n$, where d denotes the distance from the transmitter to the receiver node and the falloff exponent n was chosen to be four. It was shown in [9] that the Minimum Energy Wireless Network design applies to any propagation exponent $n \geq 2$. We do not address fading in this paper and assume diversity techniques are used to combat its effects.

The simulation results were collected for three different traffic densities given by the set $\lambda = \{1, 0.1, 0.01\}$ where λ denotes the packet inter-arrival time in seconds. The number of bits per packet is set at 2240 bits. The average data rate for each node is thus set at $R_{avg} = 2240/\lambda$ bits/second and is the same for all nodes.

In each simulation run, the broadcast period, T_B, is set to 5 seconds, which has been specifically chosen to accommodate the largest possible time that may be needed for the enclosure computation. The cost distribution period, T_C, is set to 0.5 seconds as this provided enough time for the cost tables to converge. After the cost distribution period, the data communications period starts. This period is chosen such that on average approximately 100 packets are generated per node during the simulation.

A. *Blockage Versus Maximum Data Rate*

We plotted the percentage of blockage in the network versus the maximum data rate (also called the "cutoff data rate") that a node can accommodate. If data rates are above this cutoff, the incoming packets are dropped as necessary. We graph the domain of cutoff data rates, which corresponds to the range of blockage percentages from 45% down to 0% for each network topology and for a given traffic density. The plot for $\lambda = 1$ can be seen in Figure 3; the other values of λ have similar plots.

As expected, the cutoff data rate necessary for zero percent blockage is lower for lighter traffic densities. When we compare the blockage percentages for the three different topologies for a fixed traffic density, we see that the blockage percentage for the topology of three clusters of five nodes in each cluster has a sharper increase in blockage percentage with respect to decreasing cutoff data rates. This increase is not as sharp in the case of five clusters of three nodes in each cluster. We explain this by noting that as we increase the number of nodes in a cluster while decreasing the cutoff data rate per node, the nodes providing the connectivity between clusters experience an increase in the packet-dropping rate approximately proportional to the number of nodes in the cluster. These results give

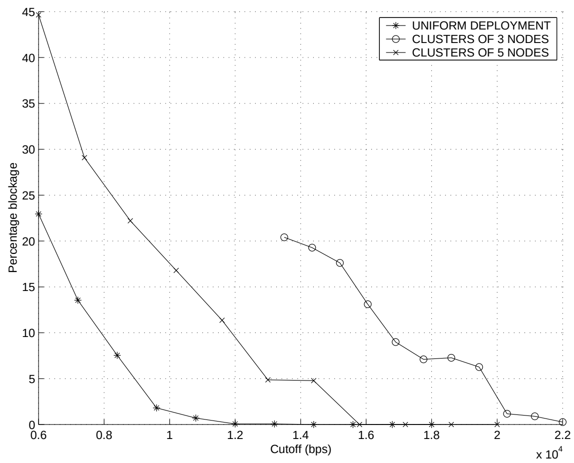

Fig. 3. Blockage Percentage for $\lambda = 1$

a perspective on the limitations each node bandwidth imposes on the network. Conversely, node bandwidths can be allocated according to a target blockage probability.

B. Broadcast Period Energy Consumption

In this section, we plot a histogram of the observed energy consumption during the broadcast period, T_B, of our network. The energy consumption is measured as the energy expenditure of each node during the transmission of packets within the broadcast period. The broadcast period energy consumption is a good indicator of the energy penalty incurred during the neighbor search period of the Minimum Energy Wireless Network. The results for the uniform deployment and small cluster topologies can be seen in Figures 4 and 5, respectively.

The broadcast energy histogram for the topology with uniform deployment shows that most broadcast energies are in a narrow range between 0.9×10^{-5}Joules and 1.5×10^{-5} Joules. This is intuitive since uniformly deployed nodes have similar enclosure radii. As we look at clustered deployments, the average enclosure radii increase since nodes have to search farther for neighbors. Hence, it takes more broadcast energy for these nodes to find their enclosures. This fact is reflected in Figure 5 for a small cluster deployment . These results give a sense of the average power consumption and its relation to different topologies.

C. Power Consumption

We plot the histogram of the power consumption per node during the data transmission period, T_D, for a traffic density of $\lambda = 0.1$ in Figures 7 and 6. We note that the power consumption per node scales appropriately for higher traffic densities and thus can be omitted from discussion. In addition, the results for uniformly deployed nodes and medium-sized clusters were seen to be similar, allowing the discussion to focus on the presented figures.

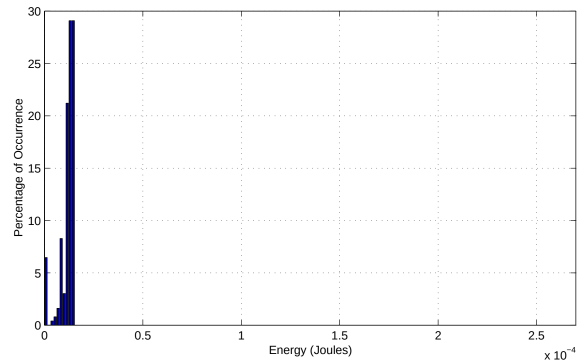

Fig. 4. Broadcast Energy for Uniform Deployment

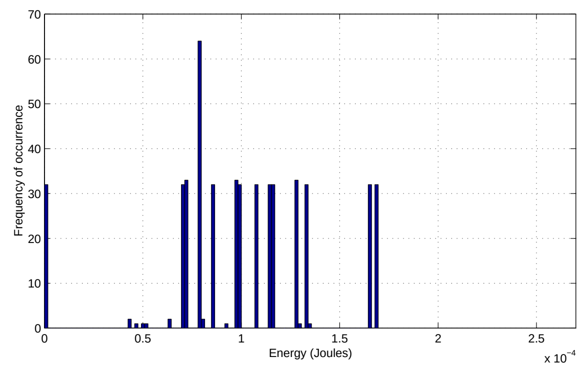

Fig. 5. Broadcast Energy for Small Clusters

These graphs clearly display the trade-off between topologies with and without clusters. In the case of deployments in clusters, about half or more of the nodes in the cluster use the remaining half of the nodes as relays to other clusters. Consequently, these nodes have to transmit to only within a very small area around themselves and appear as the large peak around zero Watts in the histograms. The nodes through which transmissions to other clusters are relayed and have larger enclosure radii experience larger power consumption as indicated by the interspersed power consumption values to the right of the histograms.

This phenomenon is not observed for uniform deployment since the enclosure radius for each node is about the same. The variations in the uniform deployment case are accounted for by the central nodes that carry most of the relayed traffic.

D. Time to Enclosure

An important parameter in the performance of the Minimum Energy Wireless Network is the maximum time needed for all nodes in the network to be enclosed. Not only does the time to enclosure place a lower bound on the duration of the broadcast period but also indicates of the delay penalty incurred in setting up the network. The enclosure time is reported by each node upon enclosure, the results of two deployments can be seen

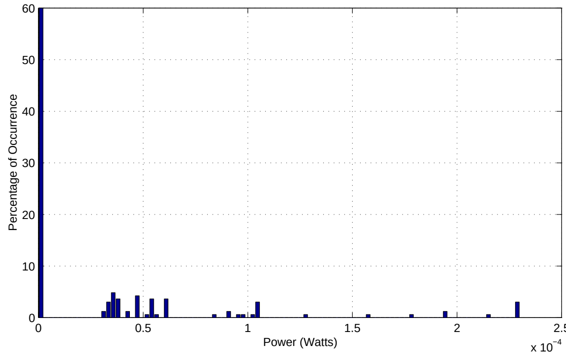

Fig. 6. Communication Power for Small Clusters

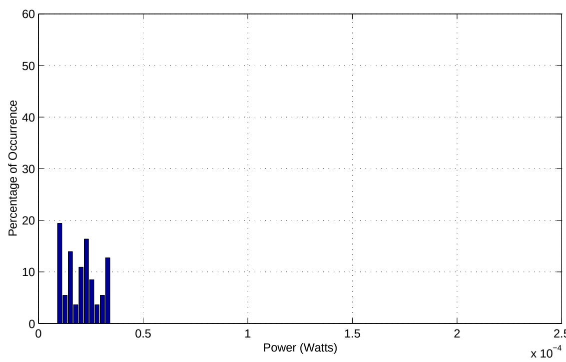

Fig. 7. Communication Power for Uniform Deployment

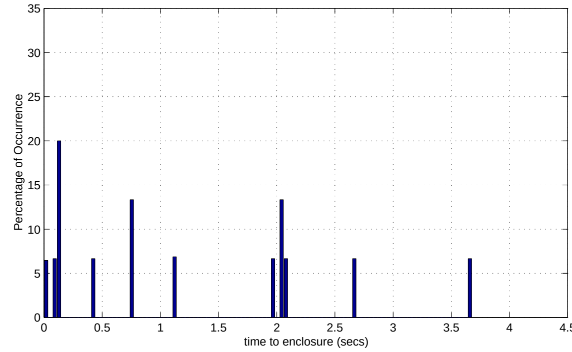

Fig. 8. Enclosure Time for Small Clusters

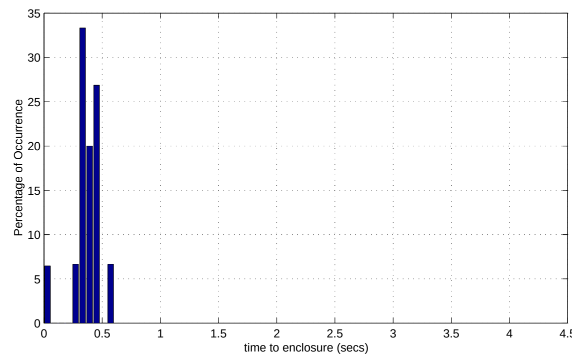

Fig. 9. Enclosure Time for Uniform Deployment

in Figures 9 and 8. Due to similarities between the results of the medium-sized and small-sized cluster deployments, only uniform and small-sized cluster deployments are plotted.

For the deployment with five clusters of nodes (with three nodes in each cluster), one or two nodes in each cluster are enclosed within a short period whereas the other ones are enclosed only after having found their neighbors in the other clusters. In contrast, for uniform deployment, the enclosure times fall within the narrow range of 0.3 to 0.45 seconds, which is accounted for by the uniformity of the inter-nodal distances.

IV. Conclusions

In this work, we studied the power, throughput, and latency (time to enclosure) performance of the Minimum Energy Wireless Network via an OPNET simulation. We also presented an acknowledgement scheme for node identification in such an ad-hoc network. We observed that this multi-hop communications network is well suited for low to medium data rate applications with power-constrained portable devices.

The simulation results display the strength of the Minimum Energy Wireless Network in adjusting to different types of deployment (uniform, medium-size clusters and small clusters). In contrast to other existing protocols, which use cluster heads and hence do not offer much robustness, the Minimum Energy Wireless Network is equally well suited for clustered and uniform deployment, as relay nodes implicitly act as cluster heads only under certain topologies. These advantages suggest the Minimum Energy Wireless Network is ideal for the proposed ad-hoc network environment.

References

[1] Z. J. Haas, "The Routing Algorithm for the Reconfigurable Wireless Networks", *ICUPC '97*, San Diego, CA, October 12-16, 1997

[2] Z.J. Haas, "On the Performance of a Medium Access Control Scheme for the Reconfigurable Wireless Networks", *Sixth WINLAB Workshop*, Kluwer Academic Publishers, 1997

[3] Z.J. Haas, "The Relaying Capability of the Reconfigurable Wireless Networks", *VTC '97*, Phoenix, AZ, 1997

[4] W.C. Jakes, *Microwave Mobile Communications*, IEEE Press 1994

[5] David B. Johnson, "Routing in Ad Hoc Networks of Mobile Hosts", *Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications*, pp. 158-163, IEEE Computer Society, Santa Cruz, CA, December 1994

[6] L. Kleinrock and J. Silvester, "On the behavior of multi-hop packet radio networks", *Proceedings of the IEEE*, vol. 75, No. 1, pp. 156-167, January 1987

[7] Barry M. Leiner, Donald L. Neilson, Fouad A. Tobagi, "Issues in Packet Radio Network Design", *Proceedings of the IEEE*, vol. 75, No. 1, pp. 6-20, January 1987

[8] Theodore S. Rappaport, *Wireless Communications: Principles and Practice*, Prentice Hall, p. 69-122, 1996

[9] Volkan Rodoplu, Teresa H. Meng, "Minimum Mobile Wireless Networks", *Journal of Selected Areas in Communications*, August, 1999

[10] John M. Rulnick and Nicholas Bambos, "Mobile power management for wireless communication networks", *Wireless Networks*, pp. 3-14, vol. 3, No. 1, March 1997