In Proc. of First International Workshop on Sensor Networks and Systems for Pervasive Computing (PerSeNS 2005) , Kauai Island, Havaii, USA, March 2005

When Does Opportunistic Routing Make Sense?

Rahul C. Shah
Univ. of California, Berkeley
rcshah@eecs.berkeley.edu

Sven Wiethölter
Technical Univ., Berlin
wiethoel@tkn.tu-berlin.de

Adam Wolisz
Technical Univ., Berlin
wolisz@tkn.tu-berlin.de

Jan M. Rabaey
Univ. of California, Berkeley
jan@eecs.berkeley.edu

Abstract

Different opportunistic routing protocols have been proposed recently for routing in sensor networks. These protocols exploit the redundancy among nodes by using a node that is available for routing at the time of packet transmission. This mitigates the effect of varying channel conditions and duty cycling of nodes that make static selection of routes not viable. However, there is a downside as each hop may provide extremely small progress towards the destination or the signaling overhead for selecting the forwarding node may be too large. In this paper, we provide a systematic performance evaluation, taking into account different node densities, channel qualities and traffic rates to identify the cases when opportunistic routing makes sense. The metrics we use are power consumption at the nodes, average delay suffered by packets and goodput of the protocol. Our baseline for comparison is geographic routing with nodes being duty cycled to conserve energy. The paper also identifies optimal operation points for opportunistic routing that minimizes the power consumption at nodes.

1. Introduction

The traditional approach to designing routing protocols for sensor networks has been to decouple the routing layer from the MAC layer. However, several real-world constraints suggest that this design philosophy is fundamentally flawed - even while attempting to optimize the quality of paths, the routing protocol can potentially generate *poor quality* paths which have low availability, low reliability and high energy consumption. The primary reason for this is that the quality of a wireless channel fluctuates significantly with time [18, 13] and the routing protocol is unaware of this fluctuatio when choosing the next hop. Additionally, nodes aggressively duty cycle to reduce their energy consumption, decreasing their availability for routing.

To alleviate many of these shortcomings, the concept of *opportunistic routing* for dense sensor networks has been proposed ([19, 20, 1, 12]). Opportunistic routing is based on geographic routing which is predicated on every node being aware of its neighbors and their specific locations. In geographic routing, the network layer of a node selects a next hop forwarder to be the node that is furthest towards the destination. This information is then sent down to the MAC layer which waits till it can achieve rendezvous with the selected node. However, in sensor networks, availability of nodes can be disrupted significantly, hence the MAC layer may suffer a significant delay and energy overhead in retransmitting the packet till it can complete the transmission successfully.

Opportunistic routing extends the idea of geographic routing, by using some node that is awake and available for routing at the time the packet needs to be transmitted. The way it works is by integrating the network layer and MAC layers so that the network layer passes down a set of candidate forwarders and the MAC layer takes a final decision on the node to use depending on current connectivity. It still uses the node location information to inform this decision, however, the particular choice of forwarding node depends on the policies of the specific protocol variant. The different varieties of opportunistic routing, however, have fairly similar performance characteristics as demonstrated later. Obviously, this approach would be attractive in dense networks where the number of potential forwarders is large.

The concept of opportunistic routing is very powerful and well-suited to sensor networks where there are significant disruptions to node availability. However, there is a cost associated with the MAC layer trying to ascertain node connectivity at the time a packet needs to be transmitted, which may negate its advantages. In fact, the overhead can be expected to grow in dense networks (with high neighborhood cardinality) and might potentially be even higher than the cost incurred due to repeated transmissions in geographic routing. Hence we would like to compare the performance of opportunistic routing with geographic routing and determine the conditions under which one approach makes more sense than the other.

This paper deals with investigating that question by conducting a systematic performance evaluation. It explores the performance of opportunistic routing for different node densities, channel quality and traffic rates and compares it to geographic routing. Transition points are found using simulations for each of the above parameters to identify when one approach provides more benefits than the other. A second contribution is to identify optimal operation points for various scenarios for opportunistic routing, i.e. the duty cycle of nodes that minimizes the power consumption.

The rest of the paper is organized as follows. Section 2 discusses some background work in this area followed by a short description of the specific protocols compared in this paper in Section 3. The simulation setup is described in Section 4 and the results are presented in Section 5. Finally, the paper concludes with Section 6.

2. Background

Many variants of geographic routing protocols ([6, 5, 7]) have been proposed as efficient ways to deal with changes in network topology. The distinguishing characteristic of this class of protocols is the use of node location information to route packets geographically towards the destination by forwarding to a neighbor that is located furthest towards the destination node. Note that in sensor networks, location information is necessary anyway since sensor data is almost always useless if it is not accompanied by location information. Hence using geographic routing protocols does not impose any additional overhead of a locationing algorithm.

In contrast to geographic routing, a number of recently proposed protocols fall under the category of *opportunistic routing protocols*. These protocols are all based on geographic routing, but dynamically choose the forwarding node based on node availability. The different variants proposed differ in the policies and selection procedure of the forwarding node, though they have fairly similar performance as shown in Section 5.

For e.g., Biswas and Morris in [1] propose a variant where the sender node transmits the packet with a specific priority of receivers specified in the packet. A system of slotted acknowledgements follows which informs all the nodes of the highest priority node that received the packet successfully. This node then forwards the packet ahead, while the other nodes drop the packet. Thus this scheme chooses the best placed node currently available. Another scheme that was proposed was in [3] where the idea of anycasting at the MAC layer was introduced which is similar in spirit to opportunistic routing. However, none of the above two works had detailed simulation results or analysis to show the efficacy of opportunistic routing.

One of the most detailed works in this area is GeRaF ([19, 20]) where the authors define a protocol that chooses the furthest node towards the destination among all the nodes that are closer to the destination than the current node. For this, they define sets of priority regions with nodes closest to the destination getting highest priority. Further, they define a fairly complicated handshake system to minimize collisions among nodes within a priority region. This will be shown later to be unnecessary since the network power consumption is minimized at extremely low wakeup rates for nodes, such that ≤ 1 node is awake on average for forwarding.

Opportunistic routing takes into account the duty cycling of nodes, hence another set of related works are topology management schemes. SPAN [2], STEM [11] and GAF [17] are all different approaches to that problem. SPAN identifies multiple sets of disjoint sets where each set provides connectivity to the whole network. On the other hand, STEM allows nodes to sleep periodically and uses beacons to rendezvous with the targeted node. Finally, GAF defines square grids where all nodes in neighboring grids can communicate with each other. All nodes within a grid are equivalent from routing purposes, hence nodes can share the routing load of the grid and sleep the rest of the time.

Multiple MAC rendezvous schemes have also been proposed in the literature. For e.g., [9] proposed two types of pseudo-asynchronous schemes - transmitter initiated (TICER) and receiver initiated (RICER). WiseMAC [4] is another interesting pseudo-asynchronous rendezvous scheme that uses preamble sampling to rendezvous with nodes while trying to minimize the power consumption.

3. Description of compared protocols

3.1. Routing protocol

The basic protocol used for comparison is geographic routing. Since we did not want to complicate the results by different versions of geographic routing, we just considered the common part of all geographic protocols - greedy forwarding. This just involves forwarding the packet to the neighbor closest to the destination. If the packet gets *stuck* due to no node being available for forwarding, it is dropped. Hence we do not consider any mechanism to route around voids or obstacles since different geographic protocols (that differ in their obstacle avoidance mechanisms) have varying power, delay and goodput performance.

For opportunistic routing, we used a specific variant called region-based opportunistic routing which was discussed in [12]. In the protocol, the network layer chooses

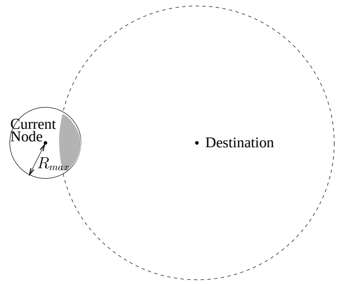

Figure 1. Lens shape as forwarding region

a forwarding region that consists of all nodes closer to the destination than the current node. Hence this forwarding region is geographically defined as a lens formed by the radio range of the current node (assumed to be circular for simplicity) and the circle centered at the destination node with the radius equal to the distance between the destination and the current node (Fig. 1). This forwarding region information is passed down to the MAC layer that forwards the packet to a node that is available for routing at that time. Similar to geographic routing, packets that get stuck and do not have any node closer to the destination are dropped. Also, the performance of this protocol will be shown in Section 5 to be similar to that proposed in [19], hence we can consider it as being representative of opportunistic routing in general.

3.2. MAC protocol

We use TICER ([9]) with two channels - data (only for data packets) and control (for control packets such as RTS, CTS and ACK packets) channels - as the MAC protocol when using geographic routing. However, when opportunistic routing is used, the MAC protocol is a variant of TICER. The difference is that we augment TICER so that it can handle two kinds of addresses from the network layer - single nodes or a region (lens) address. If the MAC layer receives an address of a single neighboring node from the network layer (as in geographic routing), it behaves in exactly the same fashion. However, the behavior changes somewhat for region addresses.

The difference lies in the fact that when a region RTS packet is sent out, it addresses several nodes located in a region rather than a single node. To avoid collisions between replying nodes, a short randomized, nonpersistent CSMA is introduced. All nodes that receive this region RTS beacon start a randomized sensing period. The CTS is only sent out in case the medium was idle for the whole sensing duration. If nodes detect the channel to be busy, they suppress their own CTS as they assume that another node transmitted a CTS already. Therefore the sending node has to wait for the maximum sensing time + one CTS duration between two consecutive RTSs. If it receives a CTS correctly, it transmits the data packet to the forwarding node. On successful reception of the packet, the forwarding node completes the transaction by transmitting an acknowledgement.

Table 1. Simulation Parameters

Geographic RTS, CTS, ACK	72 bits
Opportunistic RTS	88 bits
MAC service data unit (MSDU)	200 bits
Bit rate	40 kbps
$f_{carrier}$	1.9 GHz
Receiver Sensitivity	-80 dBm
P_{sleep}	40μW
$P_{receive}$	2.5 mW
$P_{transmit}$	4.5 mW
$P_{rx\_tx\_turnaround}$	2 mW
Radiated power	1 mW
$\alpha_{pathloss}$	3.5
Radio range	10 meters

4. Simulation Setup

4.1. Simulation Model

The simulation study was carried out using the discrete event simulator OMNeT++ [16] enhanced by the TKN Wireless Framework [14]. Varying number of nodes were randomly placed in a $50m \times 50m$ grid. Poisson distributed traffic was generated at nodes on the edge of the network with destinations being the opposite edges of the grid so that the amount of traffic seen at all points in the network is reasonably constant. In addition, a circular radio range model was assumed, while the interference range was about 1.5 times the radio range (the interference vs. radio range depended on the radio parameters in [10]).

Finally, the bit error model in [8] is mapped to a packet error model. Since [8] showed that the run length distribution of bit errors is heavy-tailed, a Pareto distribution can be used to approximate the distribution and to obtain bit errors in a packet. Hence the shape parameter α of the Pareto distribution determines the quality of the channel with higher values of α signifying worsening channels. Various simulation parameters are specified in Table 1.

4.2. Metrics and Evaluation

As the power consumption is the most critical metric in sensor networks, we use the mean power per node averaged over all nodes as our first metric. The end-to-end packet latency is the second metric and is used to consider the tradeoff between power savings and larger delays due to lower duty cycles. The third metric is the goodput which is the fraction of received to generated packets. For all simulations, we used the stochastic evaluation tool Akaroa [15] to ensure statistical significance for our results with a confidence level of 95% and a precision of the mean value equal to 5%.

Note that we will mostly present our results in terms of the node wakeup rate rather than duty cycle. The reason is that the duty cycle varies with the traffic load, length of packets, number of retransmissions required etc., hence it is not an independent variable we can control exactly during a simulation run.

5. Simulation Results

5.1. Varying node density

To see the performance of the two routing schemes for different node densities, we varied the density from 6 to 20 neighbors per node. The traffic generation rate at each of the edge nodes was a packet every 5 seconds. For a particular node density, the power consumption changes as the wakeup rates of the nodes are varied. Hence the point corresponding to the lowest power consumption was taken to represent the power performance for each node density. The channel was kept fixed at medium quality ($\alpha = 0.33$ for the Pareto distribution). The results are shown in Fig. 2. Opportunistic routing performed better than geographic routing in all cases (since nodes have to wait for a shorter time before finding a suitable next hop forwarder, thus expending less energy), though the 10% improvement was less than we expected. The primary reason was the high fraction of energy consumed at a node for periodically waking up and listening for any packets that may need to be routed (idle monitoring power), which was the same for both protocols.

Intuitively it can be seen that opportunistic routing reduces the progress per hop compared to geographical routing, increasing the total number of hops. At the same time, the choice of multiple nodes should reduce the amount of per-hop delay. This tradeoff is clearly demonstrated in Fig. 3 where the end-to-end delay performance of opportunistic routing is better than geographic routing only after the node density is > 9 neighbors per node. It is also interesting to note that the latency of opportunistic routing remained relatively constant with node density while the latency of geographic routing kept increasing. The reason for the increase in latency for geographic routing is due to the decrease in the optimal wakeup rate (for min. power) with changing node density. Since the inter-wakeup times were changed in increments of 0.05 seconds, that accounted for the step-like increase in latency in the figure (the optimal wakeup rate remains constant for $2 - 3$ different node densities before changing). On the other hand, in the case of

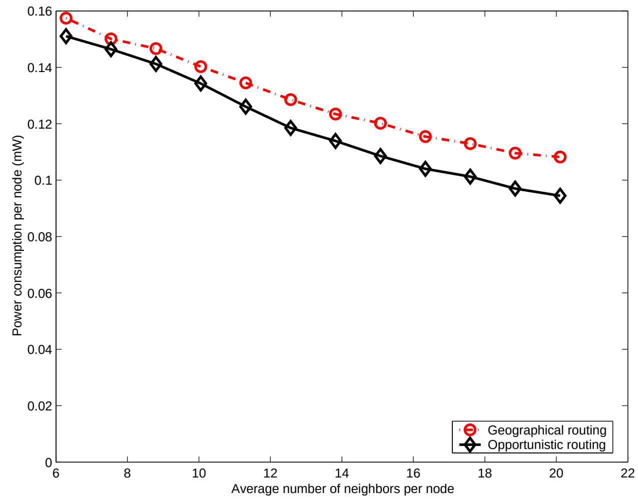

Figure 2. Power consumed per node for geographic and opportunistic routing as the density of nodes changes. Optimal wakeup rates are assumed for each density.

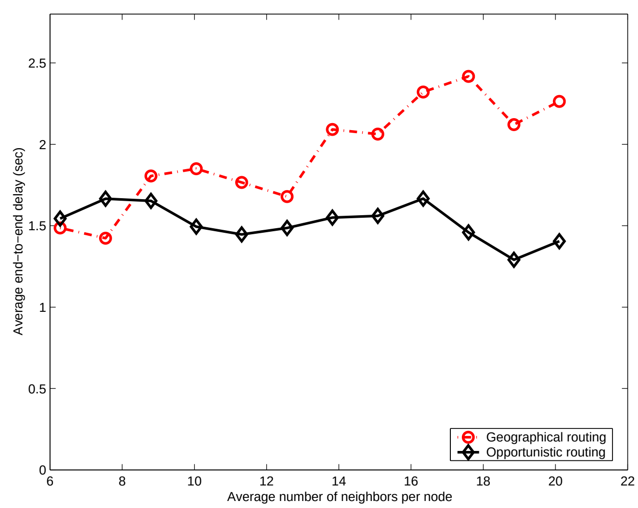

Figure 3. End-to-end delay for geographic and opportunistic routing as the density of nodes changes. Minimum power wakeup rates are assumed for each density.

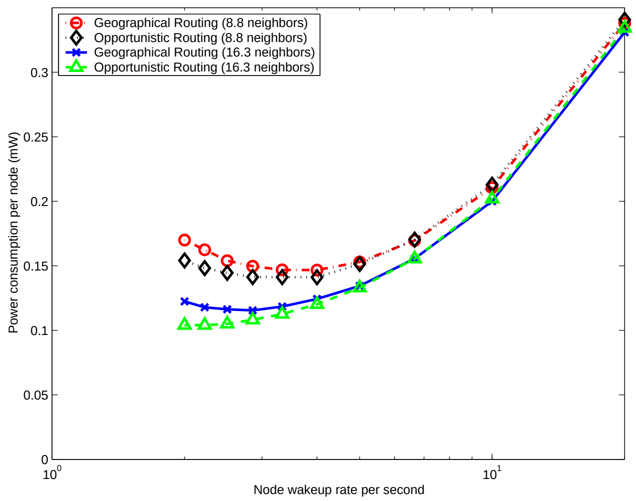

Figure 4. Power consumed per node as the wakeup rates are varied for two different node densities.

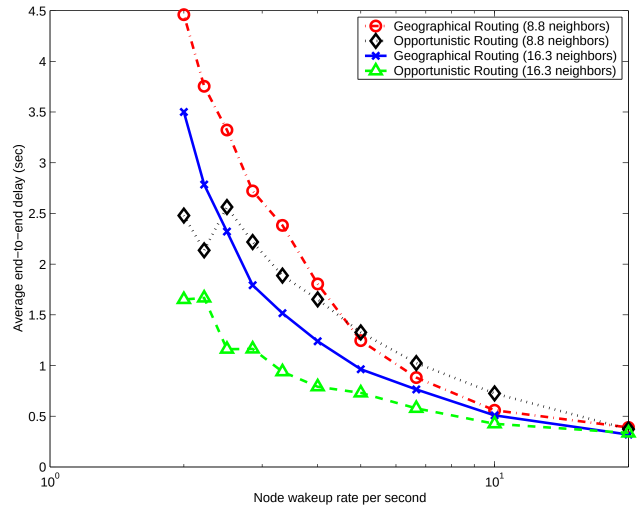

Figure 5. End-to-end delay as the wakeup rates are varied for two different node densities.

opportunistic routing, the increase in number of forwarding nodes counteracted the decrease in wakeup rates, leading to a relatively constant delay curve. Hence, the power and delay plots put together show that the benefits of opportunistic routing kick in only after a density of $\approx 9 - 10$ neighbors per node and increase steadily as the node density increases.

To see the effect of wakeup rate on the power and delay performance, Figs. 4 and 5 plot those two metrics for geographical and opportunistic routing as the wakeup rate is changed for two different node densities of 8.8 and 16.3 neighbors per node. The power consumption does not change dramatically around the optimal wakeup rate, which is good since it would be difficult to maintain the optimal wakeup rate in a real deployment where things like neighborhood sizes, traffic etc. would keep changing. The optimum wakeup rates and corresponding duty cycles per node as the density of nodes changes can be seen in Table 2. The table clearly shows the optimum duty cycle of nodes is such that ≤ 1 node is awake in the forwarding region on average.

It is also useful to compare the performance with GeRaF ([19]) to see if the results are valid for the entire class of opportunistic routing protocols. Using Eq. 15 of [19], the minimum power consumption of GeRaF can be calculated to be $0.1511mW$ and $0.1101mW$ for the two node densities of 8.8 and 16.3 neighbors per node (compare with $0.141mW$ and $0.104mW$ for opportunistic routing). Hence, as mentioned in Section 3, the different variants of opportunistic routing have similar performance. However, the more complex signaling protocol in GeRaF results in a slight power wastage and is not really required for such low duty cycles.

Table 2. Optimum wakeup rates and corresponding duty cycles per node for different node densities

Avg. number of neighbors	Wakeup rate/sec	Duty cycle
6.3	3.3	1.6%
8.8	3.3	1.6%
11.3	2.8	1.4%
13.8	2.2	1.1%
16.3	2.2	1.1%
18.8	2.0	1.0%

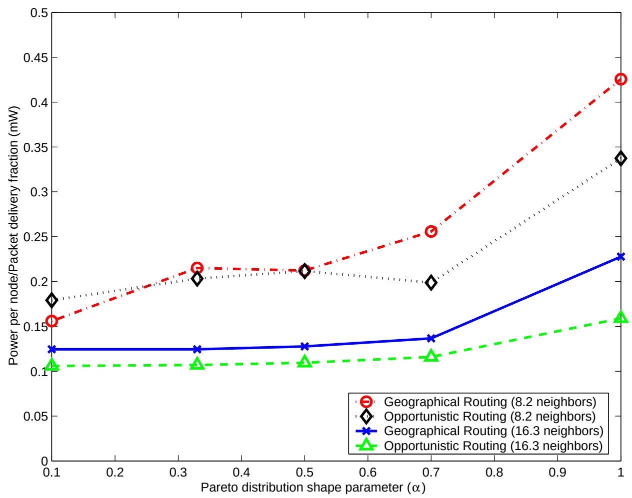

Figure 6. Power consumed per node divided by the goodput as the channel quality is varied. Higher value of α means a worse channel.

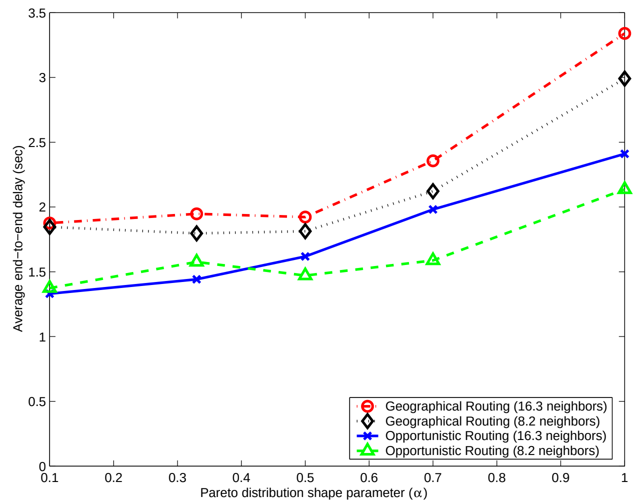

Figure 7. End-to-end delay as the channel quality is varied. Higher value of α means a worse channel.

5.2. Varying channel quality

To see the effect of channel quality on the performance of opportunistic and geographical routing, the value of α, the shape parameter of the Pareto distribution, was varied from 0.1 to 1.0 which is a very good channel to a very bad channel, respectively.

As shown in Fig. 6, for the lower node density of 8.2 neighbors, geographical routing actually outperformed opportunistic routing for a very good channel and was about the same for a medium channel in terms of the power consumption. However, for the higher node density of 16.3 neighbors per node, opportunistic routing was always better, though the gain increased as the channel was worse. Note that the y-axis plots the power consumed per node divided by the goodput. The reason for this is that while the goodput was about the same for geographical and opportunistic routing for good to medium channels, there was a significant difference in the goodput for bad channels. For such channels, geographical routing had a goodput about $10-15\%$ lower than opportunistic routing since packets were dropped if they did not get through in four attempts (one initial attempt and a maximum of three retransmissions). Hence the power was normalized for comparison purposes.

For the delay (Fig. 7), opportunistic routing performed better for all channel qualities, though the difference increased as the channel grew worse. That can be expected since the number of retransmissions would increase as the channel worsens, while the diversity of nodes in opportunistic routing would alleviate that problem to some degree.

5.3. Varying traffic rate

The final quantity to be varied was the average inter-generation time between application packets at the nodes injecting traffic into the network. Keeping the value of α for the channel at 0.33 which is a medium channel, the inter-generation time was varied from 1 second to 50 seconds. The power consumption per node is shown in Fig. 8. Since the channel was pretty good, the difference between opportunistic and geographic routing was about 10%, similar to Fig. 2. In fact, for the high traffic situation (inter-generation time of 1 second), geographic routing was about the same or better than opportunistic routing depending on the node density. The reason is that geographic routing requires much fewer hops than opportunistic routing. At the same time, for high traffic rates, the power spent in transmitting and receiving data packets becomes comparable to idle monitoring power. Hence the savings in the total number of packet transmissions start outweighing the other advantages of opportunistic routing, providing for net power savings for geographic routing.

6. Conclusions

In this paper, we tried to answer the question as to when opportunistic routing made sense over geographic routing. We used a particular variant of opportunistic routing to represent that class of protocols and used detailed simulations to compare the performance. The simulations also showed the similarity in performance between two variants of opportunistic routing. The metrics measured were the power

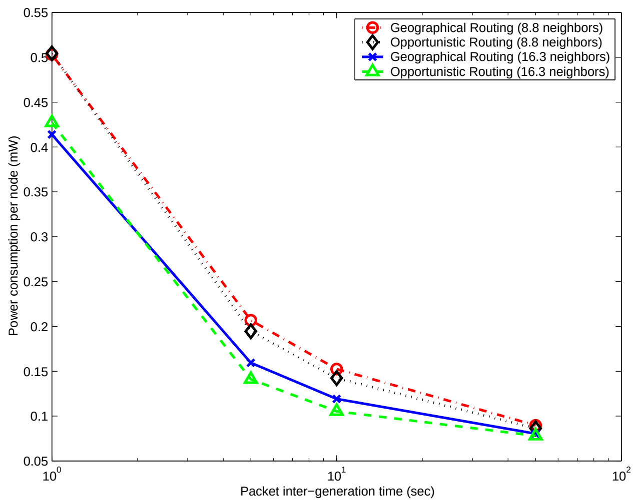

Figure 8. Power consumed per node as the traffic generation rate is varied.

consumption, end-to-end delay and goodput.

In terms of node densities for a moderate channel, it was clear that opportunistic routing was suitable only for higher node densities like more than $9-10$ neighbors per node. Even so, the gains were modest - about 10% improvement in power and 40% in delay. However, opportunistic routing was definitely superior when the channel quality worsens. On the other hand, if the channel was very good and the node density low, then it was found to be best to stick with geographical routing. Geographical routing also becomes more efficient for higher traffic scenarios.

One good thing is that the power consumption is not very sensitive to small variations in the node duty cycle around the optimal point (for both opportunistic and geographic routing). However, the biggest power consumer for both the protocols remains the idle monitoring power. Hence any mechanism to reduce this power consumption would be extremely critical in increasing the lifetime of such networks.

References

[1] S. Biswas and R.Morris. Opportunistic routing in multi hop wireless networks. In *ACM HotNets*, Aug 2003.

[2] B. Chen et. al. SPAN: An energy-efficient coordination algorithm for topology maintenance. In *IEEE/ACM MobiCom*, July 2001.

[3] R. R. Choudhury and N. H. Vaidya. Mac-layer anycasting in ad hoc networks. *ACM SIGCOMM Computer Communication Review*, 34:75–80, Jan 2004.

[4] A. El-Hoiydi et. al. WiseMAC: An ultra low power MAC protocol for the WiseNET wireless sensor network. In *ACM Sensys*, 2003.

[5] R. Jain, A. Puri, and R. Sengupta. Geographical routing using partial information for wireless ad hoc networks. *IEEE Personal Communications Magazine*, Feb 2001.

[6] B. Karp and H. T. Kung. GPSR: Greedy perimeter stateless routing for wireless networks. In *IEEE/ACM MobiCom*, Aug. 2000.

[7] Y. Ko and N. H. Vaidya. Location aided routing (LAR) in mobile ad hoc networks. In *IEEE/ACM MobiCom*, Oct. 1998.

[8] A. Köpke, A. Willig, and H. Karl. Chaotic maps as parsimonious bit error models of wireless channels. In *IEEE Infocom*, March 2003.

[9] E. Lin, A. Wolisz, and J. Rabaey. Power efficient rendezvous schemes for dense wireless sensor networks. In *IEEE ICC*, June 2004.

[10] B. P. Otis and J. M. Rabaey. A 300μW 1.9 GHz cmos oscillator utilizing micromachined resonators. *IEEE Journal on Solid State Circuits*, July 2003.

[11] C. Schurgers et. al. Optimizing sensor networks in the energy-latency-density design space. *IEEE Transactions on Mobile Computing*, 1:70–80, Jan-Mar 2002.

[12] R. C. Shah et. al. Joint optimization of a protocol stack for sensor networks. In *IEEE Milcom*, Nov. 2004.

[13] R. Stutz. Performance analysis and optimization of a 2.4 GHz multihop, wireless self-configurable network. M.S. report, Univ of California, Berkeley.

[14] Telecommunications Network Group. Framework for simulation of mobility in OMNeT++. http://www.tkn.tu-berlin.de/research/research_texte/framework.html.

[15] N. Z. University of Canterbury. Project Akaroa.

[16] A. Varga. The OMNeT++ discrete event simulation system. In *European Simulation Multiconference*, June 2001.

[17] Y. Xu, J. Heidemann, and D. Estrin. Geography-informed energy conservation for ad hoc routing. In *IEEE/ACM MobiCom*, July 2001.

[18] M. D. Yarvis et. al. Real world experiences with an interactive ad hoc sensor network. In *IEEE ICPPW*, Aug. 2002.

[19] M. Zorzi and R. R. Rao. Geographic random forwarding (GeRaF) for ad hoc and sensor networks: energy and latency performance. *IEEE Transactions on Mobile Computing*, 2(4), Oct-Dec 2003.

[20] M. Zorzi and R. R. Rao. Geographic random forwarding (GeRaF) for ad hoc and sensor networks: multihop performance. *IEEE Transactions on Mobile Computing*, 2(4), Oct-Dec 2003.