

WaveScheduling: Energy-Efficient Data Dissemination for Sensor Networks

Niki Trigoni¹ Yong Yao¹ Alan Demers¹ Johannes Gehrke¹
Rajmohan Rajaraman²

May 17, 2004

Abstract

Sensor networks are being increasingly deployed for diverse monitoring applications. Event data are collected at various sensors and sent to selected storage nodes for further in-network processing. Since sensor nodes have strong constraints on their energy usage, this data transfer needs to be energy-efficient to maximize network lifetime. In this paper, we propose a novel methodology for trading energy versus latency in sensor database systems. We propose a new protocol that carefully schedules message transmissions so as to avoid collisions at the MAC layer. Since all nodes adhere to the schedule, their radios can be off most of the time and they only wake up during well-defined time intervals. We show how routing protocols can be optimized to interact symbiotically with the scheduling decisions, resulting in significant energy savings at the cost of higher latency. We demonstrate the effectiveness of our approach by means of a thorough simulation study.

1 Introduction

Sensor networks consisting of small nodes with sensing, computation and communication capabilities are becoming ubiquitous. A powerful paradigm that has emerged recently views a sensor network as a distributed SensorDBMS and allows users to extract information by injecting declarative queries in a variant of SQL. In deploying a SensorDBMS one should consider important limitations of sensor nodes on computation, communication and power consumption. Energy is the most valuable resource for unattended battery-powered nodes. Since radio communication consumes most of the available power, SensorDBMSs need energy-efficient data-dissemination techniques in order to extend their lifetime. An important communication pattern within sensor networks is the sending of sensor readings to a designated sensor node. Let us give two examples where this pattern arises. First, consider a heterogeneous sensor network with two types of sensor nodes: many small-scale *source nodes* with low-power multi-hop communication capabilities, and a

few powerful *gateway nodes* connected to the Internet. In this setup, data flows from the sources to the gateway nodes. Our second example is motivated by resource savings through in-network processing. In-network processing algorithms coordinate data collection and processing in the network at designated nodes called *view nodes* [27, 15]. Data flows from sources to relevant view nodes for further processing.

In order to achieve energy-efficient data flows between sources and view nodes, we address several challenges intrinsic to ad hoc network communication: minimizing collisions at the MAC layer, managing radios in a power-efficient manner, and selecting energy-efficient routes. In this paper we consider data dissemination strategies that avoid collisions (and message retransmissions) at the cost of higher message latency. We carefully coordinate transmissions between nodes allowing them to turn off their radios most of the time. Since current generation radios consume nearly as much power when listening or receiving as when transmitting [10, 20, 31], the ability to turn them off when not needed yields significant energy savings. (The idle:receive:transmit ratios observed in these studies are 1:1.2:1.7 [10], 1:2:2.5 [20], and 1:1.05:1.4 [31].)

The remainder of this paper is organized as follows. We introduce our model of a sensor network in section 2. Section 3 enumerates several variants of scheduling problems and discusses their complexity. Section 4 presents our scheduling algorithm and highlights its close interaction with the routing layer. A thorough experimental evaluation of the proposed algorithm and competing approaches is presented in section 5. We discuss related work in section 6 and draw our conclusions in section 7.

2 Preliminaries

In this section, we describe our model for sensor networks and sensor data, and then briefly outline our architectural assumptions.

Sensor Networks. We consider a sensor network that consists of a large number of sensor nodes (*nodes*, for short) connected through a multi-hop wireless net-

work [26, 17]. We assume that nodes are stationary, and that all node radios have the same fixed communication range.¹ Each node is aware of its own location, and local node clocks can be reasonably well synchronized (e.g., using GPS receivers). Nodes are battery powered and thus severely energy constrained.

Sensor Data. The raw data generated at a sensor node is gathered by one or more attached physical sensors such as temperature sensors, light sensors, etc. that measure the occurrence of events (such as the appearance of an object) in their vicinity. Each sensor is a separate data source that generates records with several fields such as the id and location of the sensor that generated the reading, a time stamp, the sensor type, and the value of the reading. Conceptually, we view the data distributed throughout the sensor network as forming a distributed database system consisting of multiple tables with different types of sensor data.

Communication Patterns and View Nodes. The sensor network performs in-network processing by collecting data from multiple sensors onto a designated subset of the nodes that we call the *view nodes*. The view nodes may either store directly unprocessed sensor readings or materialize the result of more complex processing over sensor readings. The use of views in sensor networks follows a hybrid pull-push model in which relevant data is collected and *pushed* to view nodes, from where the data can be *pulled* when queries are issued.

3 Scheduling Problems And Their Complexity

A data dissemination protocol in a sensor network has two components: a *scheduling algorithm* that activates network edges such that their transmissions do not interfere with one another and a *routing algorithm* for selecting routes for individual messages. By exploring the design space of sensor scheduling and associated routing policies, we provide a graceful tradeoff between energy usage and message latency.

The particular choice of the routing and scheduling algorithms depends on the desired performance measures. Two important performance metrics are *energy consumption* and *latency*, and it is the tradeoff between these two metrics that we study in this paper. In particular, we consider the following optimization problems with respect to both energy consumption and latency metrics: (i) finding an optimal pair of routing and scheduling algorithms; (ii) finding an optimal routing algorithm for a given schedule; (iii) finding an optimal schedule for a given collection of routes.

¹Note that future generations of nodes might have variable-range radios; we leave an extension of our approach to variable-range radios for future work.

The underlying framework for our optimization problems is as follows. We assume that the sensor nodes are located on the plane and form a multihop wireless network. For simplicity, we assume that the radio range of each node is identical and equal to 1 unit; so the nodes form a *unit disk graph*: two nodes are connected by an edge if the Euclidean distance between them is at most 1. We represent the communication workload by the rate of message generation at each node i , given by r_i , and a probability distribution that gives the probability p_{ij} that a message generated at node i is destined for node j .

Energy minimization. In the *energy minimization problem*, we are given a communication workload among the sensor nodes and view servers, and our goal is to determine a data dissemination scheme that minimizes the energy consumed in delivering all messages within a bounded delay. We adopt the following model for energy consumption. Whenever a network edge is activated, the energy consumed has two components: a *fixed start-up cost* α for turning the radio on and setting up the edge for communication, and a *variable cost* β , which is the per-message transmission and reception cost. Thus, the energy consumed by an edge activation is given by $(\alpha + \beta m)$, where m is the number of messages sent during the activation².

If $\alpha = 0$, then the energy minimization problem is equivalent to finding minimum-hop paths between the source nodes and the view servers, and hence can be solved optimally in polynomial time. On the other hand, if $\beta = 0$, then the only energy consumption is in activating edges. In this case, the problem is closely related to finding a minimum Steiner tree connecting the source nodes to the view servers. The minimum Steiner tree problem, in which the goal is to determine a minimum-weight tree in a graph that connects a given set of vertices, is known to be NP-complete, even when the nodes of the underlying graph lie on a grid, a special case of unit disk graphs. We extend the NP-hardness proof to apply for arbitrary $\alpha > 0$ and $\beta > 0$.

Theorem 3.1 *For any $\alpha > 0$ and $\beta > 0$, finding an optimal routing-scheduling pair to minimize energy is NP-hard, even when there is only one view server.*

We now consider the two other parts of the design space. First, is there an efficient algorithm to compute a set of routes that minimize energy, given an activation schedule? Once an activation schedule is fixed, the fixed energy cost is already determined by the schedule; so it only remains to optimize the variable energy cost. Optimization of the variable energy cost can be achieved in

²Our model resembles the *linear model* proposed for measuring latency in communication networks [3], in which α denotes the start-up time for an edge and β is the per-message communication cost. Since the start-up time is applied once for the network as a whole, rather than the fixed cost per edge in our model, the specific characteristics of the metric in the two models are different.

polynomial time by simply routing every message along the shortest-hop path. The other partition of the design space concerns selecting an optimal activation schedule given a set of routes. In this case, we apply a reduction from the feedback vertex set problem, which is known to be NP-hard [14], to obtain that selecting an optimal activation schedule is NP-hard. We omit the proof due to space constraints.

Theorem 3.2 *For any $\alpha > 0$ and $\beta \geq 0$, given a set of source-destination routes, the problem of finding an activation schedule that minimizes energy is NP-hard.*

Latency minimization. Given a communication workload, the latency minimization problem seeks a data dissemination protocol that minimizes average latency. Again, we consider different points on the routing-scheduling design space. It is already known that minimizing latency in an ad hoc wireless network is NP-hard even for the special case where nodes exchange messages with their neighbors only [28]; the hardness of the problem stems from the difficulty of scheduling the network edges in a non-interfering manner. The reduction in [28] can be easily extended to the case of unit disk graphs by considering the 3-coloring problem [11].

Theorem 3.3 *Finding a routing-scheduling pair that minimizes latency is NP-hard. It is also NP-hard to determine an optimal activation schedule given a fixed set of routes.*

We finally consider the problem of selecting an optimal set of routes that minimize latency given a fixed activation schedule. In order to minimize latency, we need to take into the account the time that a message may spend waiting at a node before the next edge on its path is scheduled. In Section 4.1, we discuss this aspect and describe the construction of routing tables that find the minimum latency path for each message.

Our hardness results indicate that the general problem of designing an optimal data dissemination protocol, given an arbitrary sensor workload, is hard. Of course, hardness results need not necessarily imply that obtaining approximations is hard. For instance, Theorem 3.3 relies on the hardness of coloring problems. While the chromatic number for general graphs is NP-hard to approximate to even a polynomial factor, unit disk graphs admit small constant-factor approximations (e.g., see [22]). We leave the problem of designing efficient approximation algorithms to future work, and we instead concentrate on a specific part of the design space as described in the next Section.

4 Wave scheduling and routing

In this section we present *wave scheduling*, a class of periodic edge activation schedules, and study the close interaction between scheduling and routing with respect to




Figure 1: Simple wave on a 10×10 grid.




Figure 2: Pipelined wave on a 21×21 grid.

the energy and delay metric. Our scheduling mechanism is layered on top of a protocol like GAF [32], which partitions nodes into cells and periodically elects a single leader node for each nonempty cell. Nodes determine the cell that they belong to by using distributed localization techniques [2, 8]. The size of each cell is set so that a node anywhere in a cell can communicate directly with nodes in any of its four horizontal and vertical neighbor cells. This constrains the side of a cell to have length L at most $R/\sqrt{5}$, where R is the maximal transmission range of a node. The proposed wave schedules leverage the abstraction of partitioning irregularly positioned nodes into cells organized in a rectilinear grid; they focus on reducing energy consumption by coordinating inter-cell communication. For simplicity, we assume a square rectilinear grid of $N \times N$ nodes. Cell $(0, 0)$ is located at the southwest corner of the network. Cells $(i + 1, j)$, $(i, j + 1)$, $(i - 1, j)$, and $(i, j - 1)$ are the east, north, west, and south neighbors, respectively, of cell (i, j) .

4.1 Wave Scheduling

Edge activation. In our *wave schedules*, every (directed) edge of the rectilinear grid is activated periodically at well-defined communication intervals, called *send-recv*

intervals. The interval between activations is the same for all edges and is referred to as the *period*. An edge activation $A \rightarrow B$ consists of a *contention-based* and a *collision-free* period. During the contention-based period, all nodes within cell A turn on their radios in order to run the GAF protocol. If the old leader is energy-drained, a re-election protocol selects a new leader and state (routing table and message queue) is transferred to the new leader. The remaining nodes then send all readings generated since the previous GAF period to the new leader. This adapted version of the GAF protocol avoids interference caused by concurrent leader election in nearby cells. In the collision-free period leaders of A and B turn on their radios preparing for inter-cell communication. If A has no data messages to send, it sends a special *NothingToSend* (NTS) message, which allows both nodes to turn off their radios before the end of the allotted interval. The node duty cycle is thus adjusted to local traffic. In the collision-free period a data (or NTS) message is not preceded by a pair of RTS-CTS messages, but simply followed by an ACK.

SimpleWave. The intuition behind wave schedules is to coordinate message propagation in north, east, south and west phases. For instance, during the east phase, only edges of the form $(i, j) \rightarrow (i + 1, j)$ are activated sending messages along the east direction. Owing to interference, however, we cannot schedule all of the edges along the east direction. If Δ denotes the ratio of the interference range to the transmission range, then a sufficient condition for transmissions from two supernodes (i, j) and (i_1, j_1) to avoid interference is the following: $\sqrt{(i - i_1 - 1)^2 + (j - j_1 - 1)^2} \cdot L \geq \Delta \cdot R$.

In particular, two cells (i, j) and (i_1, j) can transmit simultaneously if $i - i_1 \geq \lceil \Delta \cdot R/L \rceil + 1$, which we denote by g . In the SimpleWave schedule, we schedule together edges that are g positions apart. Figure 1 illustrates the SimpleWave schedule on a 10×10 network, with cell size $L = 100m$, yielding a g of 7. The north phase starts at time 1 and lasts for 51 send-receive intervals during which every north edge is activated exactly once. The following east phase starts at time 52. In the next interval (time 53) the pattern shifts east by one cell. Only when the wave has propagated to the eighth column (time 59) it no longer interferes with node communication in the first two columns. Notice that at time 59 one can schedule four edges concurrently: $(7, 0) \rightarrow (8, 0)$, $(7, 7) \rightarrow (8, 7)$, $(0, 1) \rightarrow (1, 1)$ and $(0, 8) \rightarrow (1, 8)$.

In a *SimpleWave*, each phase takes $(N - 1) + (g - 1)g$ send-receive intervals and the entire wave period lasts for $4((N - 1) + (g - 1)g)$ intervals. This prevents the distributed deployment of the algorithm in a dynamic network: when a new cell joins (or leaves) the network, it affects the wave period and therefore the activation times of all the other supernodes. Furthermore, every node needs to know the size of the network. Another important downside of the *SimpleWave* algorithm is that it underutilizes

the capacity of the network.

PipelinedWave. This algorithm is motivated by the need for distributed and scalable schedules that make good use of network capacity. Conceptually, a network can be divided in a number of small fixed-size $(g \times g)$ squares, where all squares have exactly the same schedule. In such a network, the schedule of an edge is determined by its relative location in the square. Since all edges within the same square interfere with one another, we can schedule at most one edge at a time. The period of the resulting schedule is $4g^2$ send-receive intervals. Two edges are scheduled concurrently if they have the same direction and the sender nodes have exactly the same local coordinates within a $g \times g$ square. The *PipelinedWave* schedules a maximum number of non-interfering edges at each send-receive interval.

The *PipelinedWave* algorithm has two important properties: i) it is easily deployable in a distributed manner, since local coordination suffices for scheduling a new cell and ii) it is scalable, because node schedules are not affected by the size of the network. When a cell gets newly occupied, the associated node waits for at most one period in order to interact with its neighbors and determine its local coordinates. By overhearing the schedules of its immediate neighbors it easily determines its own schedule. When a node enters or leaves the network, the schedules of the remaining cells do not change.

A modified version of the *PipelinedWave* algorithm does not define identical schedules for each square, but shifts schedules by g positions with respect to the schedules of the four neighbor squares. More specifically, the east wave of a square is shifted g send-receive intervals earlier than the east wave of the west neighbor square, the north wave is shifted g positions earlier than the north wave of the south neighbor square etc. A snapshot of the modified *PipelinedWave* algorithm during the east phase is shown in figure 2. The new algorithm (which is the one tested in section 5) decreases the latency of message delivery at the square boundaries. Another tunable parameter in *PipelinedWave* is the number of send-receive intervals for each direction (phase) before the wave switches to another direction. Our experiments show that this parameter has no noticeable impact on the performance of the wave schedule.

Synchronization. We briefly discuss two synchronization requirements imposed by wave schedules: i) neighbor nodes must have the same notion of time regarding their communication slot and ii) nodes in the *close* neighborhood must be well synchronized so that only edges at least g positions away are scheduled simultaneously. Acknowledging that perfect time synchronization is hard to achieve, we relax the initial requirements and propose a *fault-tolerant* version of wave schedules. If the drift between two neighbor clocks does not exceed ϵ , nodes that are g positions away from each other are synchronized within $g\epsilon$. In every edge activation, we schedule the re-

ceiver to turn on the radio ϵ time units earlier than the scheduled time according to its local clock. Recently proposed synchronization protocols for sensor networks (e.g., RBS [12] and TPSN [13]) provide tight synchronization bounds (e.g., $0.02ms$ for neighbor nodes [13]) and exhibit a nice multi-hop behavior. Their performance is bound to decay for very large networks, in which case we assume that a few GPS-equipped nodes will undertake the synchronization task for the local region.

4.2 Routing

The proposed wave schedules are TDMA-based MAC protocols that assign periodic transmission slots to inter-cell. Wave schedules are general-purpose energy-efficient MAC protocols that can potentially be combined with arbitrary routing protocols. In this section we consider two important metrics for evaluating the efficiency of a routing algorithm, namely *node energy consumption* and *message propagation latency*. An interesting outcome of our study is that energy-optimal routes do not depend on the underlying wave schedule, whereas latency-optimal routes are intrinsically coupled with it.

Energy-based routing. As noted in Section 3, minimum energy routing is achieved by routing along shortest hop paths. We adopt a simple flooding approach that evaluates minimum-hop paths from all nodes in the network to a given view node. Each node in the network maintains a small in-memory routing table of size proportional to the number of view servers. For each view server, it includes a 2-bit entry giving the direction of the next hop towards the view. This simple approach works even in the presence of "holes" (empty cells), as is shown in [21]. Dynamic node failures (which manifest themselves as the appearance of new holes) can be dealt with by a local flooding phase to repair affected routes, as in AODV, or by introduction of a greedy face-routing mode as in GPSR [5, 19]. Alternatively, a node that fails to deliver a message may store it in memory until the next flooding phase that reconstructs the tree.

Delay-based routing. We propose a *delay-based* routing algorithm that, given a certain wave schedule, minimizes message latency between a pair of source and view nodes. Each node C maintains a routing table, that contains for each view V and each neighbor N a triple $\langle V, N, d \rangle$, where d is the latency of the minimum-latency path from C to V among all paths with the next-hop being N that C is presently aware of. On updating a routing entry, node C also sends the information $\langle V, N, d \rangle$ to its neighbors. On the receipt of such a message, neighbor N^* of C does the following: i) it evaluates the time dt that a message sent over $N^* \rightarrow C$ remains at C before being forwarded with the next wave via $C \rightarrow N$ towards view V ; ii) if an entry $\langle V, C, d' \rangle$ with $d' < d + dt$ exists in the routing table of N^* , then the routing message is dropped - otherwise, the routing entry is replaced by $\langle V, C, d + dt \rangle$.

When the above distributed algorithm converges, every node has determined the minimum-latency paths to each view. Routing messages can be piggy-backed on regular or NothingToSend messages as in the case of energy-based routes.

5 Experimental Evaluation

We implemented a prototype of wave scheduling in the NS-2 Network Simulator [6] and compared its performance with two other approaches: (i) an existing tree-based scheduling and routing scheme [21] and (ii) using IEEE 802.11 with different duty cycles.

Wave scheduling. We simulate a network of 20 by 20 grid cells of size $100m^2$ each. The ratio of interference to communication range is $550/250$ and the ratios between radio idle, receive and transmit power are 1:1.2:1.6. Every edge activation between two consecutive cells lasts for 200ms. A node can send about 10 packets during an edge activation given a link bandwidth of 20kbps. The receiver wakes up 30ms before the sender in order to allow for clock drift. The size of a square in a pipelined wave is set to 8 by 8 grid cells. Experiments run for 1000 seconds and the traffic workload varies from 0 to 2500 messages. The time that a message is generated is selected at random, uniformly over the simulation period. The source location of a message is randomly selected to be any of the non-empty cells, and the destination to be any of the views. Cells containing views and empty cells are randomly distributed in the network.

We first compare the behavior of the PipelinedWave schedule under two wave routing metrics: the minimum hop-count and the minimum-delay path. Figure 3 shows the average path delay, under light load, for the two metrics, i.e. the time between a generation of a message at a source and its delivery at the destination. The minimum-energy routing metric defines paths with higher delay than the minimum-delay metric and the gap increases as we increase the number of holes from 0 to 100 (25% of all cells). The energy overhead of the minimum-delay metric was observed to be negligible.

Our second experiment shows the scalability of our scheme with respect to the number of view nodes. Figure 4 shows the average observed message delay, which captures queueing delay due to traffic. We set the number of empty cells to be 0. With more view nodes, the load is better balanced across the network, the average message propagation delay is smaller and the overall capacity of the network increases. Figure 5 shows that the energy usage of the wave does not increase with the number of views, for a given number of messages.

We also examine the impact of empty cells, on the performance of wave schedules. The number of views is 10 and a randomly selected set of 0 to 80 cells are set to be holes. Figure 6 shows that the message latency increases

with the number of holes: messages wait longer in order to make a turn to bypass a hole. The capacity of the network is only 500 messages for 20% (80) holes (the message delay increases considerably after that point), whereas it rises to more than 1500 for networks without holes. Interestingly, the average energy consumption per non-empty cell (per node) increases with the number of empty cells, as shown in Figure 7. Although fewer messages are delivered per time unit, these messages follow longer paths and every node ends up routing a higher number of messages, therefore spending more energy.

In our next experiment, we vary the number of *steps* that the pipelined wave spends in each phase (or direction) resulting in different interleaved schedules. We would like to measure the message delay and node energy consumption for waves of 1, 2, 4, 8 and 16 steps. Since different interleaved schedules select different delay-optimal routes, our expectation was that they would also perform differently in terms of node energy and message delay, especially for networks with many holes. Our experiments however showed almost no difference. In order to ensure that our results are consistent irrespective of the location of views (wrt to empty cells), we run an experiment with 200 destination nodes (views) and 80 holes. As shown in figures 8 and 9 the average energy consumption per node and the average propagation delay per message are surprisingly similar for different interleaved schedules.

Tree Scheduling We compare wave scheduling with an existing tree-based scheduling and routing scheme [21]. Trees are generated as a result of a flooding mechanism initiated at each view node. Every node selects as its parent the neighbor on the shortest path to the root (view). It is therefore expected that the paths used in tree schedules are shorter than paths used in waves. Routing in a tree is trivial: each non-view node forwards every message it receives to its parent. In a tree-based schedule, we activate edges in reverse order of their distance from the root. Every tree edge is activated for 200ms seconds, as in the case of the wave.

To generalize tree scheduling to handle multiple views, we construct a collection of spanning trees, one tree rooted at each view server. An edge activation schedule can then be derived in several ways. At one extreme is a *conservative* schedule, which is simply a concatenation of schedules for the individual trees. We define a period p of repeating the activation of every tree. If we have m views, the first tree is activated at times $\{0, p, \dots\}$, the second at $\{p/m, p + p/m, \dots\}$, and so on. We assume that the interval p/m is long enough to activate all edges of a single tree, so that consecutive activations do not overlap. In Figures 10 and 11, these schedules are referred to as *Tag-Consec-Every- p* , where p is the period between two activations of the same tree. At the other extreme, we consider *aggressive* schedules that activate all trees in parallel. In our experiments, we use the name *Tag-Parall-Every- p* to refer to aggressive sched-

ules in which all trees are activated concurrently every p seconds. For instance, we activate all m trees together at times $\{0, p, 2p, \dots\}$. Figures 12 and 13 show a graceful tradeoff between energy and delay as we increase the length of period p . We note that the most energy-efficient consecutive schedule that achieves a capacity of 1000 messages has period 60 seconds. Likewise, the most energy-efficient parallel schedule that achieves a capacity of 1000 messages is activated approximately every 12 seconds. Beyond 1000 messages (per 1000 seconds), the delay for these two schedules starts increasing and it would increase without bounds had we continued to generate messages with the same rate for longer periods.

IEEE 802.11 with different duty cycles. We also study power-conserving variants of the IEEE 802.11 protocol. We vary the duty cycle of the protocol, by turning off the radio regularly and allowing communication only during 1% to 10% of the time. The performance of the resulting schemes, named *Duty-Cycle- x* , is shown in Figures 14 and 15. Routing is performed as in tree-scheduling, i.e. messages follow the shortest paths to the views. Notice that for a load of 1000 messages we can only select duty cycles greater than 8%, otherwise the traffic exceeds network capacity and the queues increase without bound. The reader can see trends in energy and delay similar to those observed in the tree-scheduling schemes.

Comparison of waves with other schemes. In order to compare different protocols we first select a given traffic load and we only consider protocols that can serve this load without exceeding capacity, which is the point at which average delay starts increasing. In fact, we compare the most energy-efficient versions of different protocols (with 10 views and 10% empty cells): for 1000 messages, we select the variants *Tag-Consec-Every-60*, *Tag-Parall-Every-12*, *Duty-Cycle-8* and the pipelined wave with step 1. Figure 16 shows that the wave protocol has the longest delay, followed by the consecutive tree schedule, the parallel tree schedule and the 802.11 (with duty cycle 8%). The reverse pattern is observed with respect to node energy consumption in Figure 17. The wave protocol is at one extreme offering the higher energy savings (better by an order of magnitude than any other scheme) at the cost of higher delay. The 802.11 protocol with duty cycle 8% is at the other extreme offering very small message delays at the cost of higher energy. The energy-delay tradeoff of the two tree scheduling algorithms is also worth observing: activating trees consecutively (as opposed to concurrently) saves energy because it avoids interference among different trees, but it incurs higher message latencies.

6 Related Work

The advent of sensor network technology has recently attracted a lot of attention to MAC and routing protocols

that are specifically tailored for energy-constrained adhoc wireless systems.

MAC protocols: IEEE 802.11 [30] is the most widely used contention-based protocol; although nodes can periodically switch to a power saving mode, in the active periods they suffer from interference and overhearing. The PAMAS MAC-level protocol turns radios off when nodes are not communicating [29], but it requires a second channel for RTS-CTS messages. PicoNet also allows nodes to turn off their radios [4]; a node wishing to communicate must stay awake listening for a broadcast message announcing its neighbor's reactivation. In S-MAC [34, 35], nodes are locally synchronized to follow a periodic listen and sleep scheme. S-MAC does not explicitly avoid contention for the medium, but reduces the period of overhearing by sending long DATA packets annotated with their lengths. NAMA and TRAMA avoid all collisions at the MAC layer by announcing the schedules of nodes in the 2-hop neighborhood and electing nodes to transmit in a given time slot. Our waves avoid schedule propagation overhead, at the expense of having fixed slots for every edge activation.

Routing algorithms: Several routing protocols for ad-hoc networks have been proposed in the literature [24, 18, 7, 25, 23]. There has also been a plethora of work on energy-aware routing [9, 29, 36] but without considering the interplay of routing and scheduling. The TinyDB Project at Berkeley investigates tree-based routing and scheduling techniques for sensor networks [21, 16]. An energy-efficient aggregation tree using data-centric reinforcement strategies is proposed in [17]. A two-tier approach for data dissemination to multiple mobile sinks is discussed in [33].

7 Conclusions and Future Work

In this paper, we have shown a class of algorithms that allows us to trade energy versus delay for data dissemination in sensor networks. Our approach is based on carefully *scheduling* the sensor nodes such that each node can stay idle most of the time, and only turns on its radio at scheduled intervals during its turn to either receive or send a message. Our experiments show that the proposed wave scheduling algorithm results in significant energy savings at modest increases in latency.

References

[1] ACM SIGMOBILE. *Proceedings of the 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM-98)*. ACM Press, 1998.

[2] Paramvir Bahl and Venkata N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *INFOCOM (2)*, pages 775–784, 2000.

[3] B. Beauquier, S. Perennes, and O. Delmas. Tight bounds for broadcasting in the linear cost model. *Journal of Interconnection Networks*, 2(2):175–188, 2001.

[4] F. Bennett, D. Clarke, J. Evans, A. Hopper, A. Jones, and D. Leask. Piconet: Embedded Mobile Networking. *IEEE Personal Communications*, 4(5):8–15, October 1997.

[5] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *Wireless Networks*, 7(6):609–616, 2001.

[6] Lee Breslau, Deborah Estrin, Kevin Fall, Sally Floyd, John Heidemann, Ahmed Helmy, Polly Huang, Steven McCanne, Kannan Varadhan, Ya Xu, and Haobo Yu. Advances in network simulation. *IEEE Computer*, 33(5):59–67, May 2000.

[7] Josh Broch, David A. Maltz, David B. Johnson, Yih-Chun Hu, and Jorjeta Jetcheva. A performance comparison of multi-hop wireless ad hoc network routing protocols. [1], pages 85–97.

[8] N. Bulusu, J. Heidemann, and D. Estrin. Gps-less low cost outdoor localization for very small devices, 2000.

[9] Jae-Hwan Chang and Leandros Tassiulas. Energy conserving routing in wireless ad-hoc networks. In *Proceedings of the 2000 IEEE Computer and Communications Societies Conference on Computer Communications (INFOCOM-00)*, pages 22–31, Los Alamitos, March 26–30 2000. IEEE.

[10] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris. Span: A energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. *ACM Wireless Networks*, 8(5), September 2002.

[11] B. Clark, C. Colbourn, and D. Johnson. Unit disk graphs. *Discrete Mathematics*, 86:165–177, 1990.

[12] Jeremy Elson, Lewis Girod, and Deborah Estrin. Fine-grained network time synchronization using reference broadcasts. *SIGOPS Oper. Syst. Rev.*, 36(S1):147–163, 2002.

[13] Saurabh Ganeriwal, Ram Kumar, and Mani B. Srivastava. Timing-sync protocol for sensor networks. In *Proceedings of the first international conference on Embedded networked sensor systems*, pages 138–149. ACM Press, 2003.

[14] M. Garey and D. Johnson. *Computers and intractability: A guide to the theory of np-completeness*. 1979.

[15] Abhishek Ghose, Jens Grossklags, and John Chuang. Resilient data-centric storage in wireless ad-hoc sensor networks. In *Proceedings of the 4th International Conference on Mobile Data Management MDM 2003*, pages 45–62, 2003.

[16] Joseph M. Hellerstein, Wei Hong, Samuel Madden, and Kyle Stanek. Beyond average: Towards sophisticated sensing with queries. In *2nd International Workshop on Information Processing in Sensor Networks (IPSN '03)*, page to appear, 2003.

[17] Chalermek Intanagonwiwat, Ramesh Govindan, and Deborah Estrin. Directed diffusion: A scalable and robust communication paradigm for sensor networks. pages 56–67. ACM SIGMOBILE, ACM Press, 2000.

[18] David B Johnson and David A Maltz. Dynamic source routing in ad hoc wireless networks. In Imielinski and Korth, editors, *Mobile Computing*, volume 353 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, 1996.

[19] Brad Karp and H. T. Kung. GPSR: greedy perimeter stateless routing for wireless networks. In *Mobile Computing and Networking*, pages 243–254, 2000.

[20] O. Kasteln. Energy consumption. Technical report, ETH-Zurich, 2001.

[21] Samuel R. Madden, Michael J. Franklin, Joseph M. Hellerstein, and Wei Hong. Tag: A tiny aggregation service for ad-hoc sensor networks. In *OSDI*, 2002.

[22] M. Marathe, H. Breu, H. Hunt III, S. S. Ravi, and D. Rosenkrantz. Simple heuristics for unit disk graphs. *Networks*, 25:59–68, 1995.

[23] V. Park and S. Corson. Temporally-ordered routing algorithm (tora) version 1 functional specification. Internet Draft, <http://www.ietf.org/internet-drafts/draft-ietf-manet-tora-spec-02.txt>, 1999.

[24] Charles Perkins and Pravin Bhagwat. Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers. In *ACM SIGCOMM'94 Conference on Communications Architectures, Protocols and Applications*, pages 234–244, August 1994.

[25] Charles E. Perkins. Ad hoc on demand distance vector (aodv) routing. Internet Draft, <http://www.ietf.org/internet-drafts/draft-ietf-manet-aodv-04.txt>, October 1999.

[26] G. J. Pottie and W. J. Kaiser. Embedding the Internet: wireless integrated network sensors. *Communications of the ACM*, 43(5):51–51, May 2000.

- [27] Sylvia Ratnasamy, Brad Karp, Li Yin, Fang Yu, Deborah Estrin, Ramesh Govindan, and Scott Shenker. Ght: A geographic hash table for data-centric storage. In *First ACM International Workshop on Wireless Sensor Networks and Applications (WSNA)*, 2002.
- [28] A. Sen and M. Huson. A new model for scheduling packet radio networks. In *Proceedings of IEEE Infocom*, pages 1116–1124, 1996.
- [29] Suresh Singh, Mike Woo, and C. S. Raghavendra. Power-aware routing in mobile ad hoc networks. [1], pages 181–190.
- [30] IEEE Computer Society. Wireless LAN medium access control (mac) and physical layer specification. IEEE Std 802.11, 1999.
- [31] M. Stemm and R. Katz. Measuring and reducing energy consumption of network interfaces in hand-held devices. *IEICE Transactions on Communications*, E80-B:1125–1131, 1997.
- [32] Ya Xu, John Heidemann, and Deborah Estrin. Geography-informed energy conservation for ad hoc routing. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 70–84, 2001.
- [33] Fan Ye, Haiyun Luo, Jerry Cheng, Songwu Lu, and Lixia Zhang. A two-tier data dissemination model for large-scale wireless sensor networks. In *Proceedings of the Eighth Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2002.
- [34] Wei Ye, John Heidemann, and Deborah Estrin. An energy-efficient MAC protocol for wireless sensor networks. In *Proceedings of the IEEE Infocom*, pages 1567–1576, 2002.
- [35] Wei Ye, John Heidemann, and Deborah Estrin. Medium access control with coordinated, adaptive sleeping for wireless sensor networks. Technical Report ISI-TR-567, USC/Information Sciences Institute, January 2003.
- [36] Yan Yu, Ramesh Govindan, and Deborah Estrin. Geographical and energy aware routing: A recursive data dissemination protocol for wireless sensor networks. Technical Report UCLA/CSD-TR-01-0023, University of Southern California, May 2001.




Figure 3: Delay vs energy routing




Figure 4: Effect of views on delay




Figure 5: Effect of views on energy




Figure 6: Effect of holes on delay




Figure 7: Effect of holes on energy




Figure 8: Effect of steps on delay




Figure 9: Effect of steps on energy




Figure 10: Delay: consecutive trees




Figure 11: Energy: consecutive trees




Figure 12: Delay: parallel trees




Figure 13: Energy: parallel trees




Figure 14: Delay: 802.11




Figure 15: Energy: 802.11




Figure 16: Comparing schemes




Figure 17: Comparing schemes