Analysis, Comparison, and Optimization of Routing Protocols for Energy Harvesting Wireless Sensor Networks

David Hasenfratz, Andreas Meier, Clemens Moser, Jian-Jia Chen, and Lothar Thiele
Computer Engineering and Networks Laboratory (TIK)
Swiss Federal Institute of Technology (ETH) Zurich, Switzerland
Email: hdavid@ee.ethz.ch, {a.meier, cmoser, jchen, thiele}@tik.ee.ethz.ch

***Abstract*—Energy harvesting has been steadily gaining interest in the wireless sensor network community. Instead of minimizing the energy consumption and maximizing a network's operational time, the main challenge in energy harvesting sensor networks is to maximize the utility of the application subject to the harvested energy. One major challenge is to maximize the data delivery rates by exploiting the spatial variations of environmental energy. While there exists a multiplicity of energy-aware routing protocols for sensor networks without energy harvesting capabilities, only a small number of routing protocols have been published which explicitly account for energy harvesting. In this paper, we analyze and compare three state-of-the-art routing algorithms. While the original algorithms assume an idealized medium access control (MAC), a lossless wireless channel and global knowledge, we show that these assumptions lead to delusive results. We detail these finding by showing the influenc of a low-power MAC protocol, a realistic wireless channel and the protocol overhead. Moreover, we show how to optimize the parameters of the MAC protocol for a given network configuration By conducting various evaluations, we identify that our modifie version of the R-MPRT algorithm outperforms the evaluated algorithms in scenarios where little energy is harvested from the environment.**

Keywords: Routing Protocols, Energy Harvesting Systems, Wireless Sensor Networks.

I. INTRODUCTION

Wireless sensor network (WSN) systems have been widely adopted for monitoring physical or environmental properties over a large area, e.g., temperature, pressure, notion, luminosity and vibration. Examples are WSNs for volcano monitoring [19], habitat monitoring [18] and permafrost monitoring [3]. In general, a WSN consists of several wireless sensor nodes that send the sensed data to a common base station for further data processing. As most sensor networks are deployed in stringent surroundings, the nodes and the system should be self-powered, self-reconfigured and self-healing. Since the sensor network is usually wide spread and the communication of limited range, multi-hop routing is required.

To prolong the lifetime of wireless sensor networks, power management under given performance requirements has become one of the most important aspects in the system. No matter how carefully one designs the system, nodes eventually run out of energy, if the only source for energy is a battery cell. The advance of energy harvesting circuits has enabled the possibility to convert and store solar energy efficientl . Recently, the emerging technology of energy harvesting has provided a means for sustainable embedded systems, which have a theoretically unlimited operation time. For example, Heliomote [7] and Prometheus [8] are two of the firs prototype sensor nodes with solar panels to harvest energy from the sun. The harvested energy can be simply used to recharge the primary energy source (e.g., a battery). Hence the system provides two sources of energy: (1) the initial battery when the system is deployed, and (2) the continuous but not necessarily constant energy harvesting rate.

Energy management for WSN has been widely explored in the literature to reduce the energy consumption for packet routing, e.g., [1], [12], [13], [16], focusing on classical battery-operated networks. Here, the objective is to maximize the network lifetime under a given workload. This is accomplished by distributing the workload as uniformly as possible over the whole network. Recently several approaches were proposed [9], [17] that give the nodes a theoretically unlimited lifetime by using solar cells. Since the initial battery charge will be depleted eventually, power management strategies have to be revised in order to decide how to use the harvested energy effectively on the long-run. As a result, routing protocols have to consider the harvested energy to decide how to route packets to the base station. This requires a paradigm shift in the design of routing algorithms.

Concerning energy awareness for sensor nodes with environmental energy supply, Zeng et al. [21] propose a routing algorithm that takes into account distance information and the link qualities. Their algorithm requires that all nodes know about the position, which usually not available in sensor network deployments. More general, Bogliolo et al. [4], Lin et al. [14], [15], and Lattanzi et al. [11] have started to explore how to maximize the workload under given environmental constraints. That is, nodes with higher energy harvesting rates (e.g., nodes that are exposed to more sunlight) are more preferable for relaying packets to the base station than those with lower rates. Specificall , the *Randomized Max-Flow* (R-MF) [4] algorithm uses an extended version of the Ford-Fulkerson algorithm to calculate the maximum fl w from the source to the destination. The probability of sending a packet over an edge is proportional to the calculated maximum fl w through that edge. The *Energy-opportunistic Weighted Minimum Energy* (E-WME) [14] algorithm annotates each node with a cost. Thereupon the shortest path from the source to the destination is calculated with respect to this node cost. The *Randomized Minimum Path Recovery Time* (R-MPRT) [11] algorithm assigns a cost to each edge and calculates the shortest path with respect to this edge cost.

The afore mentioned routing algorithms for energy-harvesting systems are theoretically sound, but the comparison metrics used are very diverse and none of them can be used to compare the routing algorithms among each other. It is difficul to see when an approach can benefi or how an approach outperforms others. Moreover, these approaches all assume an ideal medium access control (MAC) and perfect wireless links, which might lead to an illusory performance. Furthermore, these algorithms rely on global network information, such as the complete topology or the up-to-date shortest path (with respect to a define cost). The overhead for obtaining this global information is not taken into account.

In this paper we make the following contributions in the area of routing algorithms for energy-harvesting systems:

We provide the firs realistic comparison of routing algorithms for energy harvesting systems. We contribute a comparison framework that is based on the WSN simulator Castalia. This framework allows to detail the different effects when transitioning from a idealized environment to a real one, that (1) uses a real low-power MAC protocol, (2) uses a real wireless channel, and (3) generates protocol overhead for disseminating cost information. We show that the MAC protocol has the greatest influenc on the data delivery rate and further detail that this delivery rate greatly depends on the parameterization of the MAC protocol. For the performance metric we show that the average and worst-case packet loss are very well suited. Finally we propose a novel cost metric for the R-MPRT routing protocol and show that this simple protocol outperforms all other protocols considered.

The remainder of this paper is organized as follows. Section II presents the models considering the network, the energy and the workload. Section III discusses medium access control in WSNs and introduces a class of WSN MAC protocols that is especially suited for energy harvesting systems. Section IV introduces the three routing algorithms evaluated in this paper and presents our modificatio for the R-MPRT algorithm, which all are evaluated in Section V. Finally, Section VI concludes the paper.

II. System Models

In the following, we introduce our system model of an energy harvesting wireless sensor network. For this, we adopt common modeling assumptions that have been used in related literature, too. In particular, we present our models of the network structure, the energy generation/consumption and the workload within a sensor network.

A. Network Model

The network consists of one single base station and any number of sensors and routers. Sensors are able to sense a physical quantity, create data packets as well as receive and transmit them. Routers cannot sense the environment and do only forward data packets. All packets in the network are routed in a multi-hop fashion to the base station. This means that the path from the sensor to the base station consists of one or multiple edges.

The WSN is represented as a directed graph $G = (V, E)$. The vertices $v \in V$ represent the nodes (i.e., the sensors, routers and the base station). An edge $\langle u, v\rangle \in E$ represents a wireless link between the two nodes $u, v \in V$, which allows them to exchange packets.

B. Energy Model

Several technologies have been discussed for harvesting energy from a node's physical environment, e.g., solar, thermal, kinetic and vibration energy. Moreover, several prototypes have been presented which demonstrate both feasibility and usefulness of sensors nodes that are powered by solar energy [9] or vibrational energy [2].

We assume that sensors and routers are equipped with energy harvesting devices, e.g., solar panels. The available environmental energy may vary *temporally* at a single node. At the same time, there may be *spatial* variations of the harvested energy for different nodes. In this paper, we focus on *spatial* variations of environmental energy and assume that each node u has an individual harvesting power rate $p_u > 0$. The harvested energy is stored in a storage device (e.g., a battery) and we denote the stored energy $E_{C,u}$. The maximum battery capacity is define as $E_{M,u}$. We assume the base station has access to an unlimited power source $p_u = +\infty$.

The node u has available energy $E_{C,u}$ to sustain packet processing. If the energy of the node drops below $E_{\text{disable}} \geq 0$ the node gets disabled. A disabled node can neither sense nor forward data. If the node has harvested sufficien energy, i.e., $E_{C,u} > E_{\text{enable}} > E_{\text{disable}}$, the disabled node is enabled again.

Since radio communication is the main energy consumer in most sensor networks, we assume (without loss of generality) that data sensing and packet creation consumes negligible energy. Hence, each node only spends energy e_r to receive and energy e_t to transmit packets. We assume that these energies are constants since most sensor nodes send with constant transmission power. Note, however, that our framework can be easily extended to other scenarios where, e.g., transmission power is a function of the transmission range. We denote $e^{u,v}_{routing}$ the total energy that is needed to receive a packet at node u and transmit it to a neighbor node v.

Packets are retransmitted up to R times in case of a transmission failure, which increases the energy consumption for sending a packet. Assuming a packet success rate ρ_s results in an expected number or transmissions per packet of

$$(1-\rho_s)^R(R+1) + \sum_{r=0}^{R-1}(r+1)(1-\rho_s)^r\rho_s. \qquad (1)$$

If we have an estimate of the packet success rate ρ_s, we can account for the additional costs for sending retries, by multiplying the transmission costs e_t with the expected number of transmissions.

C. Workload Model

We are using uniform monitoring as workload model. Sensors are periodically sensing a physical quantity and create a data packet of the sampled data that is sent towards the base station. All the sensors are sampling with the same period ψ. Since there is no synchronization between the sampling rates of the single sensors, the sampling times are uniformly distributed among the sensors.

The packet size of the sampled data is fi ed. We assume that no data compression or in-network data aggregation is performed, i.e., all packets have to be forwarded to the base station.

III. Medium Access Control (MAC)

As already mentioned, the MAC layer has been idealized in all previous work ([4], [11], [14], [15]) on energy harvesting WSNs. In this section, we discuss the basic design trade-offs for different classes of low-power MAC protocols and decide on the most suitable one for energy harvesting WSNs.

On common sensor networks platforms, the radio device consumes the most energy, even though special low-power radio transceiver (e.g., TI CC2420, Semtech XE1205) are being used. These radios spend energy p_{on} in the order of tens of milliampere not only for sending and receiving data but also for idle listening. Hence, having the radio on at all time depletes a single AA battery within days, even if no message is sent or received. The only way for saving energy in the required order of magnitudes is by duty cycling the radio, i.e., putting the radio to a sleep mode, results in a greatly reduced energy consumption $p_{\text{off}} \ll p_{\text{on}}$. This duty cycling basically trades off latency and bandwidth for decreased energy consumption. There are various approaches to achieve this, and dozens of WSN specifi MAC protocols have emerged in recent years. The WSN MAC protocols can be divided into ones that are based on a random access scheme and to the ones that require a (global) schedule [10].

A global arbitration (e.g., TDMA) usually requires an expensive initial arbitration and further regular synchronization messages. Especially for energy harvesting systems, where nodes might stop the communication temporarily due to the lack of energy, this initial arbitration is performed multiple times. The class of schedule-based WSN MAC protocols is therefore not very well suited for harvesting systems.

For random access MAC protocols the so called low-power-listening (LPL) scheme has been shown to be very energy efficien for low data rates [10]. Using LPL the radio is powered down most of the time and is switched on every wake-up interval of T_w time units in order to check for ongoing communication. Since the sender does not know when the receiver is waking up, a long (and expensive) preamble has to be sent for ensuring that the receiver is listening when the data packet is being sent. The problem with sending long preambles is that all neighboring nodes overhear the communication and spend energy in the same order as the receiver-sender pair, which of course greatly limits the possibility for energy-aware routing. In order to reduce the overhearing, the preamble can be replaced by a strobe preamble that consists of consecutive (very short) packets containing the receiver's address.

A packet-based LPL MAC protocol is well suited for harvesting systems, since (1) it is mainly the sender-receiver pair that spends the energy and (2) no initial arbitration is required. For our evaluation we therefore choose to use SpeckMAC [20], a member of the packet-based LPL MAC protocols. An important observation with LPL protocols is, that they can be optimized for the expected amount of traffic In particular, the energy consumption for regular wake-ups is inversely proportional to T_w, whereas the message exchanges accounts directly proportional with the data load and T_w. Hence there is an optimal wake-up interval T_w given the data load. It should be noted that the chosen T_w limits the available bandwidth to $1/T_w$ transceived packets per time unit.

IV. Routing Algorithms

Three fundamental algorithms are analyzed and compared in this paper: E-WME, R-MF, R-MPRT.

A. Randomized Max-Flow

Lattanzi et al. present in [11] an extended version of the Ford-Fulkerson algorithm that calculates the maximum fl w from the sensors to the base station. It is used to create the *Randomized Max-Flow* (R-MF) algorithm. The algorithm uses the pre-calculated maximum fl w over the edges to determine the route of a packet. The probability to route a packet over an edge is proportional to the maximum fl w through that edge.

The recovery time $t_{u,v}$ is define as the amount of time required by node u to harvest the portion of energy $e^{u,v}_{routing}$ which is needed to receive and transmit a packet along edge $\langle u, v\rangle$. This can be expressed by the division of the required energy through the harvesting power rate p_u at node u. The channel capacity of edge $\langle u, v\rangle$ is given by the inverted recovery time:

$$C_{u,v} = \frac{1}{t_{u,v}} = \frac{p_u}{e^{u,v}_{routing}} \tag{2}$$

Because all the edges which belong to the same node u have to share the same power budget, the network has to be considered as a node-constrained fl w network.

Thus it is not possible to simply use the Ford-Fulkerson algorithm [6] so solve this maximum fl w problem. In [4] an extended version of the Ford-Fulkerson algorithm is presented that is used to calculate the maximum fl w in a node-constrained fl w network. Once the maximum fl w is determined, a *static* routing table is created representing the fl w through the network.

B. Energy-opportunistic Weighted Minimum Energy

Lin et al. introduce in [14] the *Energy-opportunistic Weighted Minimum Energy* (E-WME) algorithm. The algorithm define for each node u the cost c_u which depends on the available energy $E_{C,u}$, the battery capacity $E_{M,u}$, the harvesting power rate p_u and the reception and transmission energy. The algorithm calculates the shortest path from the source to the destination with respect to this node cost. The cost c_u at node u is given as

$$c_u = \frac{E_{M,u}}{(p_u + \epsilon) \cdot log(\mu)} \cdot (\mu^{\lambda_u} - 1) \cdot e^{u,v}_{routing}, \tag{3}$$

where $e^{u,v}_{routing}$ is the energy needed to receive a packet and transmit it to the downstream neighbor node v. The two constant ϵ and μ have to be chosen appropriately. The power depletion index λ_u is define as

$$\lambda_u = \frac{E_{M,u} - E_{C,u}}{E_{M,u}}, \tag{4}$$

where $E_{C,u}$ is the available energy at node u right before processing the packet and $E_{M,u}$ is the battery capacity.

C. Randomized Minimum Path Recovery Time

Lattanzi et al. present in [11] the *Randomized Minimum Path Recovery Time* (R-MPRT) algorithm. The algorithm assigns to each edge $\langle u, v \rangle$ the cost $c_{u,v}$. The sensor calculates the shortest path to the base station, with respect to this edge cost, through all its outgoing edges. The probability of sending a packet on a path is inverse proportional to the path cost.

In the original version of the paper this cost is equal to the recovery time $t_{u,v}$ introduced in Section IV-A:

$$c_{u,v} = t_{u,v} = \frac{e^{u,v}_{routing}}{p_u} \tag{5}$$

Our simulations revealed that the algorithm performs much better if we use the available energy $E_{C,u}$ instead of the harvesting power rate p_u for calculating the cost. We implemented and evaluated both the original and the modifie version of the original R-MPRT algorithm in our simulation framework. We refer to the original algorithm as *R-MPRT-org* and to the modifie algorithm as *R-MPRT-mod*.

D. Global versus Local Knowledge

The E-WME and the R-MPRT algorithms both require up-to-date cost information from every node, in order to decide where to route the packet at every node. If global knowledge is assumed, the shortest path is determined using Dijkstra's shortest path algorithm [5]. However for a real world implementation the nodes have to learn the cost information. This can be achieved in a distributed fashion if every node calculates its cost using the neighbors cost information only (with respect to the corresponding metric). This cost, which changes over time, has to be to be announced regularly by sending beacons with a certain interval ψ_B. If a node receives a beacon, the cost metric has to be updated and a potential change has to be announced.

The R-MF algorithm requires global knowledge to create the routing tables at every node. It is not clear, how this can be implemented in a distributed fashion without generating extensive overhead in terms of memory consumption, processing and communication overhead.

V. Simulation Results

We compare the four algorithms E-WME, R-MF, R-MPRT-org and R-MPRT-mod described in Section IV. The comparison is executed on Castalia 2.2, a simulator for wireless sensor networks that is based on the OMNeT++ platform. Castalia allows to realistically model the sensor nodes and the wireless channel. The radio device is accurately modeled having different radio states with individual energy consumptions and transition times. The wireless channel uses an advanced path-loss model that is based on experiments and calculates the packet probability using a signal to interference-plus-noise ratio (SINR) model.

A. Simulation Setup

We have conducted simulations on various network topologies of different sizes. In the following, we will limit ourselves to evaluations on a 10x10 grid topology, which is illustrated in Figure 6. There is one base station, 81 sensor nodes and 18 routers that only forward packets. The base station is placed on the border of the grid, has unlimited energy available and is always listening, i.e., the radio device is always on. We evaluate the algorithms for a simulation time of 60,000 s and use 3 different seeds for modeling different behavior of the realistic channel model with its random packet loss.

As radio model the TI CC2420 transceiver is used. It is designed for low-power and low-voltage wireless applications. The energy consumption for receiving and idle listening is 62 mW, whereas the transmission costs 57 mW (0 dBm). For duty cycling the radio is put into the power down mode, decreasing the power consumption to 0.072 mW. The battery capacity is 100 joules, however the initial energy is set to 1 joule. This enables the nodes to process the firs few packets, yet requires the node to harvest the energy for processing further packets.

Using this simulation setup, we evaluate step by step the impact on the four routing algorithms going from an idealized to a realistic scenario. First we simulate with a perfect MAC protocol and an ideal channel characteristic. We then introduce a state-of-the-art low-power MAC protocol and evaluate the data delivery rate for different MAC protocol parameterization. We then analyze the impact of having a realistic model of the wireless channel and further account for protocol overhead.

We finis the evaluation by examining different harvesting rate distributions.

B. Evaluation Metrics

We defin new metrics to compare the presented routing algorithms within our simulation framework. For the metrics, we opted for the average as well as for the worst-case packet loss. They are define as follows:

- **Average Packet Loss** denotes the total number of lost packets divided by the total number of generated packets by all sensors.
- **Worst-Case Packet Loss** is the highest number of lost packets originating from a single sensor (which got lost on their way to the base station) divided by total number of packets generated by the sensor.

Having a high worst-case packet loss ratio means that there are single sensors in the network whose data rarely arrive at the base station. In other words, the worst-case packet loss reflect how "balanced" the data delivery works and may indicate blind spots which are unobserved by the sensor network. The existence of blind spots contradicts the idea of uniform sampling in the sensor network.

A packet is lost if (a) a packet cannot be created because the node is disabled (b) a packet cannot be forwarded by a sensor/router since there is not enough energy to receive or transmit the packet (c) there are collisions on the MAC layer (d) there is a transmission error due to a poor link quality. Note that packets are retransmitted $R = 3$ times, which is enough to handle packet losses due to collision and poor links, but is not enough if the receiving node is lacking environmental energy on the long run.

C. Idealized Environment

We start the evaluation having an idealized environment as assumed in [4], [11], [14]. In particular, we assume an idealized channel model without packet loss. Furthermore, we assume a perfect low-power MAC protocol that only spends energy for sending and receiving data, and uses global knowledge available from the simulation for synchronizing the sender-receiver pair. It should be noted that such a perfect synchronization cannot be realized, yet is often assumed when designing routing protocols.

Figure 1 illustrates the packet losses for a uniform energy harvesting scenario, i.e., all nodes have the harvesting power rate $p_u = 0.95$ mW. In general, the packet losses increase for shorter packet injection periods. The average packet loss of the E-WME and both R-MPRT-mod algorithms are very similar, R-MF performs slightly worse. However, already in this plot we see a bad worst-case performance of original R-MPRT-org algorithm: The sustainable packet injection rate where almost no packet losses occur is reduced by 33 % compared to the other algorithms. As already discussed, the worst-case packet loss can be very important for the decision which routing protocol to choose.

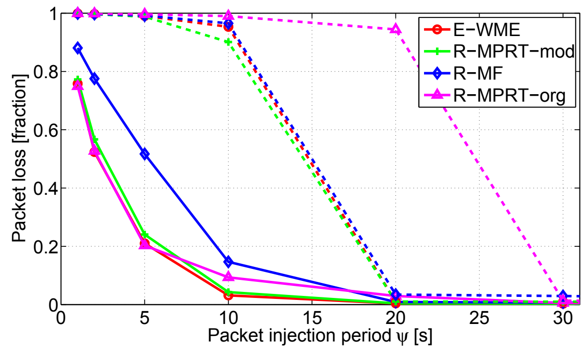

Fig. 1. Average (solid lines) and worst-case (dashed lines) packet loss of the four routing algorithms in an idealized environment with a perfect MAC protocol and an ideal wireless channel.

D. Low-Power MAC Protocol

As a firs step towards a more realistic network model, a realistic MAC protocol is used, which has to coordinate packet transmission and reception between senders and receivers. Nodes are put in a power-saving sleep mode most of the time and only wake up for short communication intervals. The base station on the other hand has its radio always on, due to its unlimited power supply. This enables the base station's immediate neighbors to send messages without the long preamble and therefore to save large amount of energy when sending messages. In the firs part of this section, we will illuminate the inherent difficultie for optimizing the MAC protocol. The second part is dedicated to a comparison of the four routing protocols taking into account the MAC layer. We assume a wireless channel with the possibility for packet loss due to collisions, but assume a perfect link quality $\rho_s = 1$.

1) MAC Parameter Optimization: For LPL-based MAC protocols, the energy spent at a node highly depends on the wake-up interval T_w. A short T_w results in spending a lot of energy for idle listening, yet makes the message transmission inexpensive. For a long T_w little energy is spent idle listening but a message exchange is getting expensive. Hence there is a traffi dependent T_w that minimizes the energy consumption. We will now investigate this trade-off for SpeckMAC, a typical LPL protocol for wireless sensor networks.

Figure 2 displays the average packet loss in dependency of the wake-up interval T_w and the packet injection period ψ. The simulation is performed for the E-WME routing algorithm; the qualitative results for the other algorithms are similar. It becomes obvious that an energetically sustainable workload (i.e., no packet loss) is only achieved for periods $\psi \geq 1000$ s. For periods $\psi < 1000$ s, the packet loss is not neglectable. At the same time, for each injection period ψ, there exists (at least) one wake-up interval $T_w \in [0.1, 04]$ s, which minimizes the packet loss. We also observe that a very short wake-up interval of $T_w < 0.1$ s results in too much energy being spent for idle listening and leads to nodes running out of energy

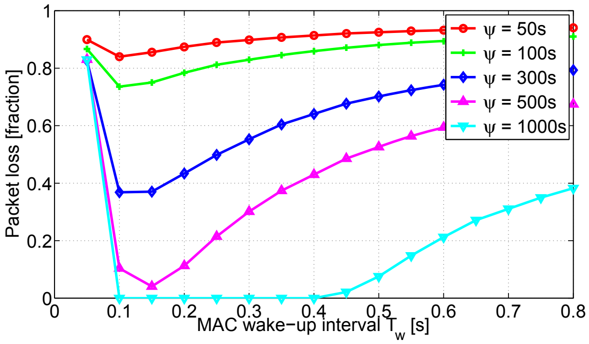

Fig. 2. Packet loss as a function of the wake-up interval T_w of the SpeckMAC protocol for different packet injection periods ψ.

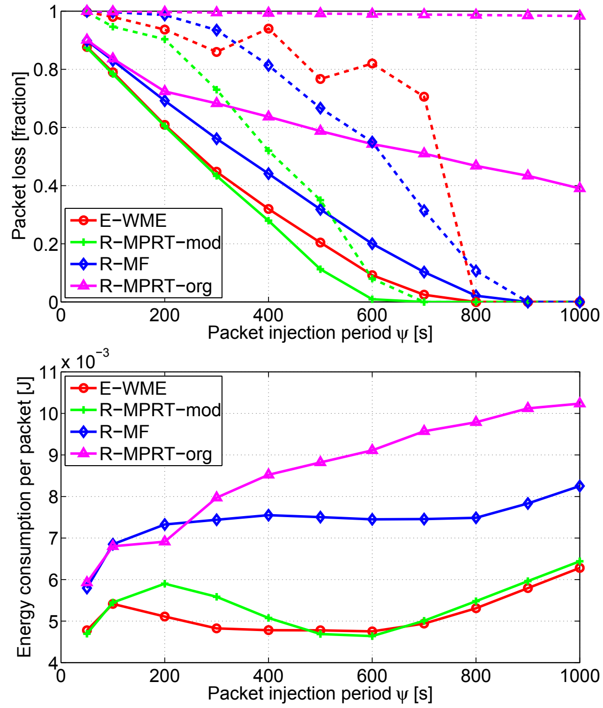

Fig. 3. Top plot: Average (solid lines) and worst-case (dashed lines) packet loss; Bottom plot: Energy consumption per successfully transmitted packet; both using SpeckMAC with wake-up interval $T_w = 0.2$ s.

frequently. On the other hand, for long wake-up intervals $T_w > 0.4$ s, too much energy is spent for transceiving messages.

In summary, using the methods presented in this section, one can determine the maximum sustainable packet injection period ψ and a suitable wake-up interval T_w for a given network configuration For the remaining analysis we use a wakeup interval of $T_w = 0.2$ s that shows to be a well balanced value for our application scenario according to Figure 2.

2) Comparison of Routing Algorithms: The top plot in Figure 3 illustrates the packet loss of the considered routing protocols. Compared to Figure 1, the general performance degradation due to MAC overhead becomes visible: Through the use of a real MAC protocol the sustainable packet injection periods are increased by a factor of 30 to 40. This fact and the more detailed results show the importance of the simulation of the MAC protocol.

Furthermore, we see huge differences between R-MPRT-mod and the other algorithms, which are not existent in the idealized environment discussed in the previous section. Having a real low-power MAC protocol, R-MPRT-mod outperforms the other algorithms with respect to the worst-case *and* the average packet loss. For packet injection periods 600s $< \psi <$ 800s the worst-case packet loss of R-MPRT-mod is even lower than the average packet loss of all other algorithms. In practice, the lowest injection periods ψ which result in a neglectable packet loss are important. The performance of R-MPRT-org is not acceptable in all our experiments and we decide to not further discuss this algorithm.

In the lower plot of Figure 3, some interesting results on the energy consumption per successfully transmitted packet are displayed. We have the minimal energy consumption for a packet injection period of 600 seconds. This is just about the energetically sustainable rate of the protocols (where almost no packet losses occur) and allows to utilize the available energy the best. Injecting more packets increases the energy consumption for the successfully transmitted packets. This can be attributed to the fact that energy is spent for messages that eventually do not reach the base station. For very short packet injection periods ($\psi < 100$ s), the energy consumption per successfully transmitted packet decreases again. The reason for this surprising effect is that most of the transmitted packets stem from one-hop neighbors of the base station that can send the packets without a long preamble.

E. Realistic Wireless Channel

Wireless sensor networks have to deal with packet loss due to the wireless channel. Instead of assuming perfect links, we used Castalia's realistic channel model with its random packet loss. That is, the simulator calculates for every link a packet success rate ρ_s, which represents the probability that a packet is successfully transmitted. With these random link qualities, there are however also links in the network that provide a very low packet success rates (close to zero), which makes such link not usable. We therefore choose to use only links having a $\rho_s > 0.8$, assuming that we have a link estimation algorithm that provides us with updated link qualities.

We run the simulation with the same configuration with the additional constraint of realistic wireless channels. The results show a rather equal amount of performance degradation along all three algorithms, which we do not show in detail. In particular, the sustainable packet injection period is increased by about 100 s, e.g., from $\psi \sim 600$ s to $\psi \sim 700$ s for R-MPRT-mod.

For the remaining analysis of the distributed algorithms and

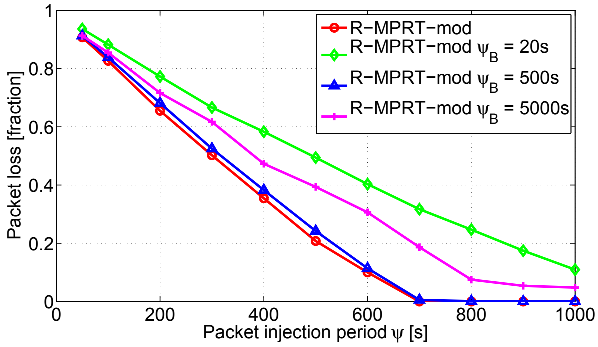

Fig. 4. Packet loss of the distributed R-MPRT-mod algorithm for different beacon intervals and the centralized one (without beacons) that assumes global knowledge.

the non-uniform harvesting rates we assume a realistic MAC and wireless channel.

F. Protocol Overhead

Until now, the routing algorithms were provided with all the required information for free. For a real-world implementation, the algorithms have to learn about their neighbors and their cost metric. This requires a distributed solution that can determine the route with information from immediate neighbors only. This can be implemented for the R-MPRT and E-WME algorithms. The R-MF algorithm on the other hand requires the information from all nodes. Since it is not clear how this global knowledge can be retrieved (especially for larger networks), the R-MF algorithm is not considered.

Retrieving information from neighboring nodes is achieved by sending beacons. In order to minimize the delay, the beacons are initiated at the base station with a certain beacon interval ψ_B. The information is then propagated in the network until all the nodes have an updated view of the current topology. The beacons create additional costs and therefore influenc the sustainability of the network. This is illustrated in Figure 4 for R-MPRT-mod only, since E-WME yields similar results. As illustrated, the distributed algorithm achieves almost the same sustainability rate than the implementation that requires global knowledge (lowest line) if an optimized beacon interval $\psi_B = 500$ s is chosen. If we choose a longer interval, we can save energy by sending fewer beacons, which however comes at the price of an outdated topology resulting in an increased packet loss. Choosing on the other hand a shorter period is not necessary, since the beacon interval of $\psi_B = 500$ s already provides sufficientl updated information and hence only wastes energy.

These results clearly show that the retrieval of the routing information cannot be ignored. For R-MF where global knowledge is required, the retrieval is not easily possible. For R-MPRT and E-WME a simple distributed version can be implemented. It is however crucial for the implementation to carefully choose an appropriate beacon interval.

G. Different Harvesting-Rate Scenarios

All simulations so far have been performed under the assumption that all nodes are exposed to the same environmental source, and thus have the same harvesting power rate p_u. To investigate the influenc of an unbalanced spatial distribution of environmental energy, we studied several scenarios and present two of them in the following:

- Harvesting Scenario 1: The harvesting power rates p_u are normally distributed. The mean value is equal to the constant harvesting rate $p_u = 0.95$ mW.
- Harvesting Scenario 2: As indicated in Figure 6, half of the nodes have a significantl lower harvesting power than the rest. Again, to make the scenarios comparable, the sum of harvesting power rates p_u is the same as in the previous scenarios.

Figure 5 shows the packet losses for the two harvesting scenarios. Although the authors of [11] claim that R-MF should approximate the maximum, optimal fl w through the network, algorithm R-MF performs very poor in this experiment, in particular for Harvesting Scenario 2. Overall, we could not validate the superior performance of algorithm R-MF in [11] in our experiments. A possible explanation for this behaviour could be the appearance of packets which are looping in the network. These packets consume a significan amount of energy. Since packets are forwarded to neighboring nodes according to probabilities, looping packets cannot be avoided by the R-MF algorithm. This effect is much stronger in non-uniform energy harvesting scenarios.

Figure 6 visualizes the packet fl ws for both the R-MF and the R-MPRT-mod algorithm for Harvesting Scenario 2. The packet fl w on an edge is indicated by the thickness of the respective line. Comparing the two plots, it becomes clear why R-MPRT-mod achieves a more balanced behavior, resulting in lower packet loss rates.

VI. Conclusion

In this paper, three recently published routing protocols for energy harvesting wireless sensor networks are analyzed and compared. Contrary to the original works, we compare the protocols under realistic scenarios, including a real MAC protocol, a lossy wireless channel and the protocol overhead. Specificall , we show how parameters of the MAC protocol can be optimized for a given harvesting scenario and network topology. Moreover, we investigate distributed versions of the algorithms that could be indeed implemented on sensor nodes. As a further contribution, we propose a modifie version of R-MPRT algorithm that outperforms its competitors significantl in terms of sustainable packet injection rates. Finally, we propose novel evaluation metrics that allow a comparison of the different algorithms.

Surprisingly, for all simulation scenarios, we found that our modifie algorithm R-MPRT-mod which uses solely the stored energy in its cost function outperforms the other algorithms.

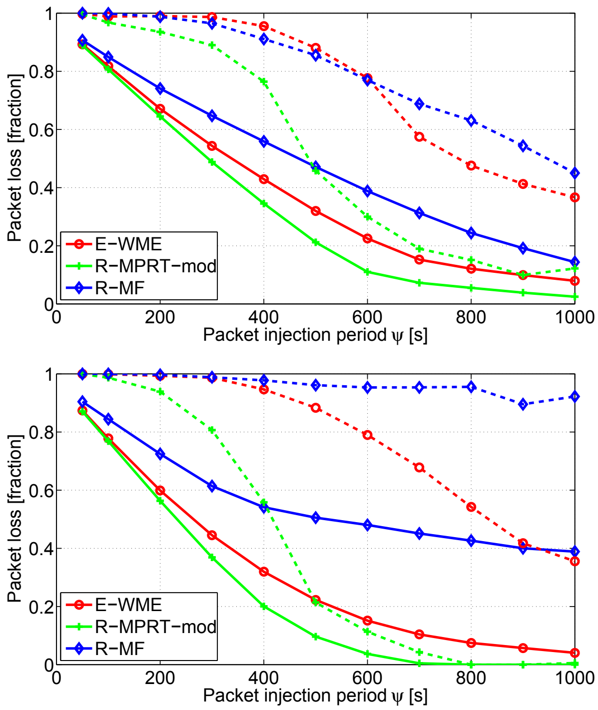

Fig. 5. Average (solid lines) and worst-case (dashed lines) packet loss for Harvesting Scenario 1 (top plot) and Harvesting Scenario 2 (bottom plot).

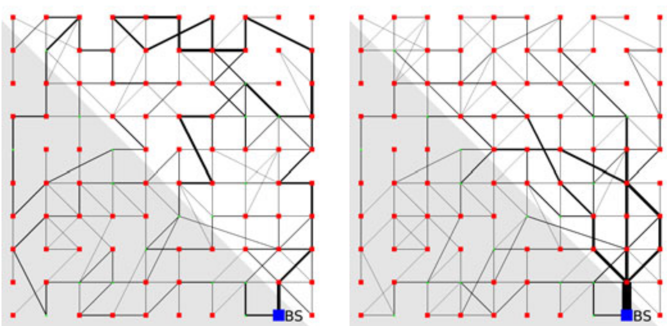

Fig. 6. Visualization of the packet fl ws on a 10x10 network topology for Harvesting Scenario 2; One base station (BS), 81 sensors (red rectangles) and 18 routers (no symbols); routing algorithm R-MF (left) versus R-MPRT-mod (right).

This comparatively simple algorithm is only indirectly aware of the harvested energy. This result is in contrast to finding in related work where a paradigm shift from "energy-aware" to "energy-harvesting aware" protocols is claimed.

Our experiments highlight that simplistic assumptions on the lower protocol layers lead to wrong results when analyzing routing protocols. In particular, it is of great importance to carefully choose the MAC protocol and parameterize it appropriately. Furthermore it must be kept in mind that wireless links are not ideal and that the knowledge of the cost metrics of neighboring nodes cannot be obtained for free.

Acknowledgements

This work is supported by the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322. This research is also supported by the European Network of Excellence on Embedded System Design ARTISTDesign.

References

[1] K. Akkaya and M. Younis. A survey of routing protocols in wireless sensor networks. *Elsevier Ad Hoc Network Journal*, 3(3), 2005.

[2] Y. Ammar, A. Buhrig, M. Marzencki, B. Charlot, S. Basrour, K. Matou, and M. Renaudin. Wireless sensor network node with asynchronous architecture and vibration harvesting micro power generator. In *sOc-EUSAI'05*, New York, USA, 2005.

[3] J. Beutel, S. Gruber, A. Hasler, R. Lim, A. Meier, C. Plessl, I. Talzi, L. Thiele, C. Tschudin, M. Woehrle, and M. Yuecel. PermaDAQ: A scientifi instrument for precision sensing and data recovery in environmental extremes. In *IPSN'09*, San Francisco, CA, USA, 2009.

[4] A. Bogliolo, E. Lattanzi, and A. Acquaviva. Energetic sustainability of environmentally powered wireless sensor networks. In *PE-WASUN'06*, New York, NY, USA, 2006.

[5] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1), 1959.

[6] L. R. Ford and D. R. Fulkerson. Flows in networks. *Princton University Press*, 1962.

[7] J. Hsu, A. Kansal, J. Friedman, V. Raghunathan, and M. Srivastava. Energy harvesting support for sensor networks. In *IPSN'05*, 2005.

[8] X. Jiang, J. Polastre, and D. E. Culler. Perpetual environmentally powered sensor networks. In *IPSN'05*, UCLA, Los Angeles, California, USA, 2005.

[9] X. Jiang, J. Polastre, and D. E. Culler. Perpetual environmentally powered sensor networks. In *Fourth International Symposium on Information Processing in Sensor Networks*, 2005.

[10] K. Langendoen. Medium access control in wireless sensor networks. In H. Wu and Y. Pan, editors, *Medium Access Control in Wireless Networks*. Nova Science Publishers, Inc., 2008.

[11] E. Lattanzi, E. Regini, A. Acquaviva, and A. Bogliolo. Energetic sustainability of routing algorithms for energy-harvesting wireless sensor networks. *Comput. Commun.*, 30(14-15), 2007.

[12] L. Li and J. Y. Halpern. Minimum-energy mobile wireless networks revisited. In *ICC'01*, 2001.

[13] Q. Li, J. Aslam, and D. Rus. Online power-aware routing in wireless ad-hoc networks. In *MobiCom'01*, 2001.

[14] L. Lin, N. B. Shroff, and R. Srikant. Asymptotically optimal energy-aware routing for multihop wireless networks with renewable energy sources. *IEEE/ACM Trans. Netw.*, 15(5), 2007.

[15] L. Lin, N. B. Shroff, and R. Srikant. Energy-aware routing in sensor networks: A large system approach. *Ad Hoc Netw.*, 5(6), 2007.

[16] V. Mhatre and C. Rosenberg. Energy and cost optimizations in wireless sensor networks: A survey. Proc. of Annual Allerton Conf. on Communication, Control and Computing, 1999.

[17] V. Raghunathan, A. Kansal, J. Hsu, J. Friedman, and M. B. Srivastava. Design considerations for solar energy harvesting wireless embedded systems. In *IPSN'05*, 2005.

[18] R. Szewczyk, J. Polastre, A. Mainwaring, and D. Culler. Lessons from a sensor network expedition. In *EWSN'04*, 2004.

[19] G. Werner-Allen, K. Lorincz, J. Johnson, J. Lees, and M. Welsh. Fidelity and yield in a volcano monitoring sensor network. In *OSDI'06*, Berkeley, CA, USA, 2006. USENIX Association.

[20] K.-J. Wong and D. K. Arvind. Speckmac: low-power decentralised mac protocols for low data rate transmissions in specknets. In *REALMAN'06*, New York, NY, USA, 2006.

[21] K. Zeng, K. Ren, W. Lou, and P. J. Moran. Energy-aware geographic routing in lossy wireless sensor networks with environmental energy supply. In *QShine'06*, New York, USA, 2006.