Generic Routing Metric and Policies for WSNs

Olga Saukh, Pedro José Marrón, Andreas Lachenmann, Matthias Gauger, Daniel Minder and Kurt Rothermel

IPVS, Universität Stuttgart
Universitätsstr. 38
D-70569 Stuttgart, Germany
{saukh, marron, lachenmann, gauger, minder, rothermel}@informatik.uni-stuttgart.de

Abstract. Energy-aware algorithms have proven to be a crucial part of sensor network applications, especially if they are required to operate for extended periods of time. Among these, efficient routing algorithms are of utter importance since their effect can be experienced by all other layers. Thus, the optimization and accurate prediction of the lifetime of the system can only be performed in the presence of accurate execution models that take energy consumption into account. In this paper, we propose a generic routing metric and associated policies that encompass most other existing metrics in the literature and use this model for the optimal construction of a routing tree to the sink. We also provide experimental results that show the benefits of using our novel metric.

1 Introduction

Sensor networks are getting more and more popular, especially in the area of intelligent monitoring, where the deployment of standard sensing equipment is too expensive or even impossible. Application domains such as habitat monitoring, environmental observation or intelligent building monitoring usually need a cheap, easily deployable and self-organized system that is able to operate for years on a single set of batteries. Such applications pose strict energy-awareness requirements to the hardware and software used.

Many classic algorithms have been analysed and optimized to make them energy-aware. A number of new algorithms have been proposed specifically for sensor networks which try to minimize energy consumption while still providing the desired QoS level. However, energy is usually not included in the model as a parameter, and when included, it is mostly in the form of a constraint rather than a real optimization parameter.

The basis of any working network, particularly a sensor network, is an efficient routing protocol. Since wireless communication is the most energy-consuming task, the routing module needs to be carefully tailored towards saving energy as much as possible. Of great interest is not the routing protocol itself, which usually gives a constant overhead for each transmitted packet, but rather the underlying routing tree topology defined by a metric, which is directly responsible for the

K. Römer, H. Karl, and F. Mattern (Eds.): EWSN 2006, LNCS 3868, pp. 99-114, 2006.

quality, stability and energy-awareness of the routing tree. Existing solutions usually concentrate on selecting best quality paths by means of end-to-end packet reception rate. In many cases energy-awareness is defined in terms of the number of retransmissions, which definitely influences energy consumption albeit only in an indirect way.

The challenge is, therefore, to find an optimal routing tree structure, which selects the best paths using QoS parameters, like reliability of a given path, and/or its energy demands. Although the compromise between these two parameters might not fit the requirements of different applications to the routing module, we expect such a generic solution to be parameterized and tunable. There might be applications that require reliable packet delivery at any price, or those that might tolerate losses but should be optimized to operate for several years. Finally, we are interested in examining the characteristics of the resulting routing tree topology defined by our metric.

The rest of the paper is organized as follows. Section 2 gives an overview of related work. In section 3 we present an example that shows the limitations of existing approaches and motivates our work. Section 4 lists our assumptions and provides a detailed description of the generic metric as well as the policies used in our model. Section 5 presents the results obtained through experimental evaluation of our model on real hardware. Finally, section 6 concludes the paper and describes future work.

2 Related Work

Relevant papers available in the literature that relate to our model can be grouped in three topic areas: routing metric definition and analysis, transmission power level tuning approaches and energy consumption modeling.

Routing Metric Definition and Analysis. Most papers in this area have been developed for traditional computer networks and for ad-hoc networks, and sometimes they also apply to sensor network research. Let us discuss three representative examples. The Shortest Path First metric (SPF) discussed in [3, 6] selects the route based only on path length information. It has been shown to be not suitable for sensor networks, since it selects the neighbor further away with the lowest link quality to route packets. An optimized version of it, called SPF(t), applies a blacklisting procedure to filter the links with quality less than t before using the SPF algorithm on the resulting topology. Clearly, SPF(t) shows better behavior than SPF, but it might lead to a disconnected routing tree, as shown in [6]. The Success Rate (SR) metric tries to find the paths with the highest end-to-end success rate as a product of link qualities p_{ij} along the path $Path$: $\prod_{(i,j)\in Path} p_{ij}$ [5]. However, it usually underestimates the path quality, since it does not take into account the possibility of packet retransmission. It might also lead to cycles in the routing graph. Finally, the ETX metric was originally developed for ad-hoc networks, but is used in sensor networks as well. Its goal is to minimize the sum of the expected number of transmissions over each link $(i, j) \in Path$: $\sum_{(i,j)\in Path} 1/p_{ij}$. Compared to the other two, this metric has been shown to behave best, especially in the presence of low mobility [1]. Other metrics based on latency, like the ones presented in [2], have shown poor behavior when used in sensor network scenarios.

Transmission Power Level Tuning. In [4], the authors have investigated means of influencing the packet reception rate of a link by tuning the transmission power level in single-hop scenarios. In [7], the dependence between the received signal strength and the packet reception rate is investigated on the physical and MAC layers. Finally, the authors of [8] used the transmission power level to regulate the number of available neighbors, and therefore, mitigate the number of collisions in dense networks.

Energy Consumption Modeling. Energy has been modeled at different levels in a number of papers. [9] presents an analytical model to predict energy consumption in saturated IEEE 802.11 single-hop ad-hoc networks under ideal channel conditions. In [10] the energy spent to send and receive a message is accurately modeled and a power aware routing protocol is proposed that adapts routes to the available power. For the approach proposed in this paper, we can use any of the low level models available in the literature, since our goal is the optimization of the routing paths themselves. [11] tries to find the optimal traffic scheduling in energy contrained networks to optimize the lifetime of the network. In [12] the authors have shown, that always using lowest energy paths may not be optimal from the point of view of network lifetime and long-term connectivity. They propose a scheme that occasionally uses sub-optimal paths to provide sustainable gains.

However, to the best of our knowledge, our work is the first to combine transmission power control, number of transmissions, energy consumption and success rate in one routing metric. In the next section we present an example which motivates the need for proper modeling and direct inclusion of energy consumption into the routing metric.

3 Motivation

Fig. 1 shows an example of a network with four nodes, where node 3 tries to send some data to the sink (node 0). There are two possible routes to reach the sink, either via node 1 or via node 2. The links $(3, 1)$, $(2, 0)$ and $(1, 0)$, $(3, 2)$ have link qualities 1.0 and 0.1 respectively. In this example, we assume that all nodes spend an equal amount of energy for each attempt to transmit a packet over any available link and that no retransmission of lost packets is in place.

The three metrics mentioned in the related work consider both paths to be equivalent. The SPF metric assigns a value of 2 to both paths because there are two hops to the sink. The SR metric guesses the end-to-end success rate to be the same $(0.1 \cdot 1.0 = 0.1)$. Finally, the ETX metric estimates the number of transmissions to be $\frac{1}{0.1} + \frac{1}{1.0} = 11$ for both paths.

However, the expected energy spent by the sensor network if the first route is selected is nearly twice as much as the energy spent in the case of the second

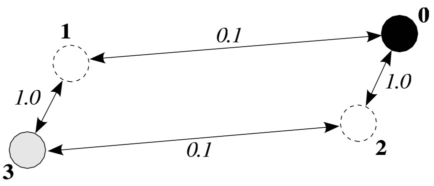

Fig. 1. Motivating example

route. The reason for this is that node 2 forwards packets from 3 only in 10% of the cases. If we assume that transmitting all packets generated by 3 costs 1 energy unit, then node 2 spends 0.1 energy units for forwarding. This way, the energy costs of the routes via node 1 are 2 energy units and via node 2 are 1.1 energy units.

Moreover, if we assume that all nodes operate at different send power levels and have different numbers of retransmission possibilities, the resulting energy demands of using either path might change considerably. Therefore, in the next section, we analyse the general case and build a model which takes these facts into account: path quality and energy consumption of the path. Based on this model, we propose a new routing metric and discuss its properties and ways to tune it by means of application-specific policies.

4 Formal Model

The goal of our approach is to determine a routing tree structure that provides the highest end-to-end packet success rate while taking the amount of consumed energy into account. For the success rate, it suffices to independently determine its value at any given point in time. All metrics in the related work section work using this method. However, including the consumed energy in the model requires the consideration of past history, since the amount of available energy depends on the history of previous system performance.

For this reason, our model contains two optimization metrics: end-to-end packet success rate, which depends mostly on the topology of the network and the environmental characteristics; and energy consumption, which, as stated above, depends on past system performance. Our approach involves the optimization of the ratio of these two metrics taking into account the following two assumptions: First, we only consider two independent parameters, transmission count and transmission power level that allow us to tune the direction of optimization. Second, we only allow these two parameters to change their values over time, and assume that they are managed by application-specific policies.

Using these values, at each point in time the metric selects the best possible route based only on the current values and not taking into account previous history. Changes to these parameters over time will be considered later when we deal with application-specific policies.

4.1 The Metric

Terminology Let $G = (N, E)$ be a fully connected undirected graph, where N is a set of nodes numbered $[0..n]$ and node 0 is called the *sink*. $\forall i, j \in N$ each *link* $(i, j) \in E$ has a pair of characteristics: $(p, e) \in (Pr, En)$.

For any two nodes $i, j \in N$, we define the *link quality* $p_{ij} \in Pr$ to be the product of the probabilities indicating that a packet sent by node i is received correctly by node j and vice versa. If no blacklisting is used, $Pr = [0, 1]$. Two nodes $i, j \in N$ are defined to be *neighbors*, if $p_{ij} \neq 0$. This implies that implicit acknowledgements can be used in our model, since neighbors can always hear each other.

The sequence of links $(i, k_1), (k_1, k_2) \dots (k_f, 0)$ is a *path* $Path_i$ from node i to the sink 0, if each link $(i, k_1), (k_1, k_2) \dots (k_f, 0)$ connects two neighbors. Without loss of generality we assume that all paths end at the sink. The metric defines a partial order over all paths. We consider only paths from the same node to the sink to be comparable. If node j is chosen to be a parent of node i along the path, then we define the notation $p_i = p_{ij}$.

Additionally, $e_i(l) \in En$ is a discrete function which represents the dependency between the transmission power level l and energy e_i needed for one transmission of a packet using this power level. En is the domain of such discrete functions that might be different for each radio type.

Each node $i \in N$ initially possesses the amount of energy Ξ_i. If possible, it might manage the transmission power level l_i of the radio and change the number of possible packet transmissions r_i. The number of possible retransmissions in the case of send failures is then $(r_i - 1)$. Both techniques have an effect on the amount of energy spent to send the packet and may influence p_{ij}. We consider l_i and r_i to be the input parameters of our model and to be tunable by application-specific policies.

Our model tries to describe and analyse the *end-to-end packet reception rate* or *path quality* referred to as $Gain_{path_i}$ and the *total energy* spent by all nodes along this path denoted as $Energy_{path_i}$. If we consider any possible path, the notation $Gain$ and $Energy$ is used.

Mathematical Description Transmitting a packet over a link can be modeled as a Bernoulli trial. However, transmitting a packet over a path is not a Bernoulli process, because if the packet is lost on one of the links, it is not forwarded any further. This fact is the basis of our model which considers *expected energy consumption* $Exp(Energy)$ and *expected path quality* $Exp(Gain)$ from the point of view of the influence of transmission power control and the available number of retransmissions.

Tuning the transmission power and applying a retransmission mechanism might look like different techniques, but both leave their print on the energy consumption of the link and its quality. By performing one retransmission, we spend twice as much energy, even though the packet might still be lost. By increasing the transmission power level we expect to increase the signal to noise ratio and hope for better link quality. However, as explained in [7], this is not guaranteed. These facts make retransmissions and transmission power control similar instruments in our model.

Let us now assume that a packet is originated at node n and should be forwarded to the sink 0 using some $Path_n$ (see Fig. 2). The link qualities along the path are $p_n, p_{n-1}, \dots p_1$ and the energy spent by each node to transmit a packet via a link in the absence of retransmissions is $e_n(l_n), e_{n-1}(l_{n-1}), \dots e_1(l_1)$. The possibility to retransmit a packet directly influences both parameters and moreover, separates the model into several cases dependent on the maximum number of times r the packet is allowed to be transmitted. Commonly, there are between 1 and 3 transmission attempts available.

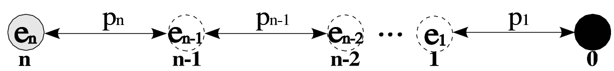

Fig. 2. Path evaluation: general case

First, we consider the case, where nodes do not have the possibility to perform retransmissions. Later, the model is extended by removing this limitation. Finally, we discuss the case when the number of transmissions is unlimited on each link.

Case 1: No retransmissions ($r_i = 1$) This is the simplest case, when the retransmission mechanism is not used and the packet is allowed to be sent only once ($r_i = 1$). We are interested in the evaluation of the path characteristics ($Energy_{path}$ and $Gain_{path}$) given the qualities p_i and power consumptions $e_i(l_i) = e_i$ of each link dependent on the current transmission power levels l_i. The mathematical expectation of each characteristic is:

$$Exp(Gain_{path})_{r_i=1} = p_n(p_{n-1}(\dots p_2(p_1)\dots)) \tag{1}$$

$$Exp(Energy_{path})_{r_i=1} = e_n + p_n(e_{n-1} + p_{n-1}(\dots + p_2(e_1)\dots)) \tag{2}$$

As can be found from the expressions, there is no need for global knowledge to calculate the values at each node. It is enough if each node propagates its accumulated values through the routing tree.

Case 2: Limited number of transmissions ($r_i = x_i \ll \infty$) Let us assume, that in this case all the nodes in the network may perform a maximum of x_i transmissions, that is, $(x_i - 1)$ retransmissions on the link $(i, i-1)$ if the packet is lost. Then the probability that the packet sent by node i is received by node $(i-1)$ is:

$$a_i = \sum_{k=1}^{x_i} p_i(1-p_i)^{k-1}$$

and node i is expected to perform

$$b_i = \sum_{k=1}^{x_i} kp_i(1-p_i)^{k-1} + x_i(1-p_i)^{x_i}$$

attempts to send a packet given $r_i = x_i$ possible transmissions. The aforementioned formula needs some explanation. After the first attempt to send a packet, it is successfully received by the node $(i-1)$ with probability p_i. However, with probability $(1-p_i)$ the packet needs to be retransmitted. In other words, with probability $(1-p_i)$ we need to perform at least two attempts to send a packet. But during the second try, only $(1-p_i)p_i$ packets are expected to be successfully delivered. Therefore, with probability $p_i(1-p_i)^{k-1}$ a packet is transmitted via link i after k transmission attempts. If we are limited to x_i transmissions, with probability $(1-p_i)^{x_i}$ we are still not able to send a packet successfully and are expected to give up. In this case we have made x_i unsuccessful attempts.

If each send attempt costs e_i energy units, node i spends b_ie_i energy units to send a packet given $r_i = x_i$ transmissions.

So, the expectation values of both metrics are:

$$Exp(Gain_{path})_{r_i=x_i} = \prod_{i=1}^{n} a_i \tag{3}$$

$$Exp(Energy_{path})_{r_i=x_i} = b_ne_n + a_n(b_{n-1}e_{n-1} + a_{n-1}(\cdots + a_2(e_1)\dots)) \tag{4}$$

Case 3: Infinite number of transmissions ($r_i = \infty$) If we have an infinite number of transmissions and $p_i \neq 0$, sooner or later the packet will be transmitted via any link $(i, i-1)$ with probability 1:

$$a_i = lim_{n\to\infty} \sum_{k=1}^{n} p_i(1-p_i)^{k-1} = 1$$

If $p_i = 0$, the link does not connect any neighboring nodes, so it cannot be selected to transmit packets. The expected number of attempts for delivery over a link is then:

$$b_i = lim_{n\to\infty} \sum_{k=1}^{n} kp_i(1-p_i)^{k-1} + n(1-p_i)^n = lim_{n\to\infty} \frac{p_i}{1-p_i} \cdot \sum_{k=1}^{n+1} k(1-p_i)^k = \frac{1}{p_i}$$

Therefore, in the case of an infinite number of transmissions, we have:

$$Exp(Gain_{path})_{r_i=\infty} = 1 \tag{5}$$

and if each send attempt at node i costs e_i energy units, we have:

$$Exp(Energy_{path})_{r_i=\infty} = \sum_{i=1}^{n} \frac{e_i}{p_i} \tag{6}$$

Discussion Analyzing already existing metrics like ETX, SR and SPF from the perspective of the model we just described, we can see, that they cover only partial cases and none takes either the possibility to control the transmission power level or to set the maximum number of transmissions into consideration. The ETX metric corresponds to equation (6) when $e_n = e_{n-1} = \cdots = e_1 = 1$ and accounts for an infinite number of transmissions, which simplifies the model but does not fit reality. The SR metric coincides with equation (1). The SPF accounts for neither energy consumption nor the path quality and in the general case is expected to behave unpredictably over both characteristics.

Therefore, the SR and ETX metrics are just two special cases of our model.

GEM Let us now use the model we have just defined to create a metric that can be optimized to find better routes. A metric is a rule which builds the routing tree and lays the foundation for its efficiency. Therefore, there are some features we would like our metric to reflect. It should be based on the model used and should be able to find a compromise between path quality (and, therefore, gain) and energy consumption. We define our metric to be:

$$\text{GEM} = \frac{Exp(Gain)}{Exp(Energy)} \tag{7}$$

GEM stands for **G**ain per **E**nergy **M**etric and, as we will see in the experimental section, behaves better in practice than other existing routing metrics. From a more theoretical point of view, GEM has some nice features:

GEM finds the best throughput paths taking energy into account. Both expected end-to-end success rate ($Exp(Gain)$) and expected resource demands ($Exp(Energy)$) are directly included into the target function and, therefore, changes to any of these parameters influence the GEM path estimation value. Moreover, GEM is tunable by changing the values of the number of transmissions r_i and transmission power levels $e_i = e_i(l_i)$ at each node, which allows for the definition of application-specific policies very easily.

The evaluation of a path as done by GEM is based on the realistic assumption that the link layer always has a limited number of retransmissions available. However, the application itself might decide to retransmit a packet as many times as needed. In this case, $1/Exp(Gain)$ reflects the expected number of transmissions needed over a path (not a link) to transmit a packet successfully at the application layer.

In order to calculate the value of GEM, there is no need to have global knowledge about the system. GEM makes the route selection based on the accumulated values for the expected end-to-end success rate and energy consumption propagated by each node starting from the sink. Additionally, it obviates the necessity of blacklisting, and, therefore, accounts for a wide range of link loss ratios and even the existence of asymmetric links.

Finally, GEM leads to the formation of a naturally structured routing tree in the sense that it decreases the probability of a decomposition of the tree into a forest. This feature comes from including energy demands directly into the metric. The nodes that are far away from the sink have a tendency to select the furthest nodes first for data forwarding. The rational behind it is that if a packet gets lost, it is better to lose it as soon as possible, as shown in the motivating example. On the other hand, the closer the node is to the sink, the more reliable and better quality paths are selected by our metric. This natural tree structure contributes to routing tree stability and, as a result, makes GEM a perfect candidate for defining hierarchical structures.

4.2 Energy-Aware Policies

Although GEM accounts for both energy consumption of a path and end-to-end path quality, it is simply a ratio that needs to be optimized. In many applications it is important to stress which of these two parameters is more important and in which way. Simple parametrization is not much of a help, since it is usually impossible for the application to define the exact relative importance between energy and gain.

Therefore, we propose the use of policies, a more complex, but application specific type of parametrization. Each policy is aimed at optimizing either the energy or the end-to-end success rate characteristic of a path and might be applicable to a certain type of applications. Therefore, a policy is responsible for setting the right values for the transmission power level and retransmission count over time, that is, taking the history of the system into account. In general, a policy should know how often the application has requested to send a packet, which signal strength it used and what number of retransmissions was needed to forward the packet to the next node along the path. This information together with data about the environment obtained through the observation of link quality changes over time and the goals of the application serve as input parameters for the policy.

Below we present some examples of policies that perform different optimizations and can be modeled using our approach. All of them are based on tuning either the transmission power level or setting an appropriate maximum transmission count for each node in the network.

Policy Examples In [8] the authors examined the *Neighborhood Based Policy* (NB), a policy that decreases the transmission power level if the number of neighbors exceeds a predefined value and vice versa. This might be useful for adaptation of routing algorithms to sparse and dense networks and is supported by our model by setting r_i to a fixed value.

In [10] the authors describe a model of the node lifetime based on the node power function ξ_i, which describes the node energy demands at each point in time. $\Psi_i(t) = \Xi_i - \int_{t_0}^{t} \xi_i(t)dt$ shows the energy left on the node at a certain point in time t[1]. The policy called *Lifetime Based Policy* (LTB) is then able to keep track of $\Psi_i(t)$, predict the lifetime of the system, and modify its characteristics by adjusting both transmission power level and transmission count.

[1] Recall that Ξ_i is the initial amount of energy on node i

Some policies might be based on tuning the maximum transmission count as well as the transmission power level. For example, it is possible to differentiate packet content and distinguish more or less important information, but still keep track of the energy spent to send packets to a sink. More important packets might be sent using a higher transmission power level or be allowed to be retransmitted more times in case of send failures (*Packet Type Based Policy* (PTB)).

Both tuning parameters can also be set individually by each node to make link qualities in the routing tree be equal. Then the goal is to select the lowest energy paths without affecting the end-to-end packet success rate (*Equal Gain Based Policy* (EGB)). In this case, GEM is expected to perform close to ETX with the only difference that the maximum transmission count is always limited, even if to a large value and, therefore, gains in energy consumption along the path will be measured. This would be the case especially in sparse and poorly structured networks.

The transmission power level as well as the number of transmissions might be used to optimize the packet reception rate of the link with the current parent (*Link Quality Based Policy* (LQB)). According to [7], increasing the transmission power level does not guarantee a better link quality. However, if LQB is based on this parameter, it should help making the links from the transitional region be more stable and, therefore, in many cases, should still increase the link quality as shown by the experiments in [4].

Obviously, these are just examples and a more exhaustive classification of policies and an analysis of their characteristics must be performed and is left for future work. However, in order to show their potential effects, we present some results of GEM combined with LQB and compare it to only using GEM.

5 Experimental Results

For the experiments presented in this paper we used 10 Tmote Sky motes based on the TI MSP430 microcontroller with a CC2420 radio module located in an office environment. The nodes were distributed in a room of 4.8 × 2.4 meters and the initial transmission power levels were set to one of the lowest possible values in oder to organize them into a multihop network.

There are some routing module settings that need to be mentioned. We have used a moving average to estimate the link quality with a window size of 10 (see [3]). The acceptable number of missed packets is equal to 5, after which the entry is deleted from the routing table. The routing table size was set to 10, so that all the nodes had equal chances to be selected as parents for packet forwarding. Initially, all nodes communicate at power level $l_i = 2$ (range is [1..31]).

We have tested the three metrics from related work and compared them in the first set of experiments with GEM without using any policy (GEM-off). In the second set of experiments we investigated the effect of the LQB policy (GEM-lqb) and compared it to GEM-off.

5.1 GEM without policies

In this set of experiments we evaluated the ETX, GEM-off, SPF and SR metrics for the case of 1, 2 and 3 maximum available transmissions and equally fixed transmission power levels. In our setting, the experiment lasted 100 seconds and was repeated a total number of five times. The graphs present the average of these five experiments.

As evaluation criteria we have included the end-to-end packet success rate, which shows the percentage of successfully received packets from each node in the network; energy consumption of the selected path from each node; hop distribution along the path from each node to the sink; route stability, that is, the average number of parent changes during one experiment; and route maintenance overhead, that is, how much information needs to be sent accross links to maintain the routing tree.

End-to-end Success Rate The three graphs on the left-hand side of Fig. 3 show the percentage of packets which were successufully delivered to the sink node by each one of the nodes. The SPF metric has the worst behavior because it selects the minimal-hop paths and, therefore, the longest low quality links to route packets to the sink. GEM is not influenced by any policy and reflects energy consumption and gain as a simple ratio. However, it shows nearly the same level of packet success rate as ETX. The reason for this is that the links from the transitional region are unstable, and, therefore, usually not considered for routing packets. This does not happen as a result of blacklisting but rather because the packets are received accidentally via such links. The SR metric shows a good packet success rate, since it is targeted at maximizing this parameter characteristic.

Energy Consumption The three graphs on the right-hand side of Fig. 3 display the energy demands of using the selected path from each node. In all our examples we calculate the energy consumption of a path in energy units. One energy unit is the amount of energy needed to send a packet at the lowest possible transmission power level ($l = 1$). Using the Tmote Sky specification it is easy to estimate how many energy units are spent if communication takes place with higher power levels.

The energy consumption of paths obviously increases with each additional retransmission. SPF has the lowest energy demands in the case of 1 transmission. However, this is not true anymore if the retransmission mechanism is available. This is because links with low qualities need additional send operations and, therefore, increase energy consumption of paths.

The SR metric shows the highest energy consumption for the paths from nearly every node. This makes sense since it is the only metric that does not take energy consumption into account. ETX does this indirectly by estimating the number of transmissions and SPF does it by minimizing the number of hops.

ETX and GEM both take energy consumption into account (ETX is just a special case of GEM with the number of transmissions approaching infinity).

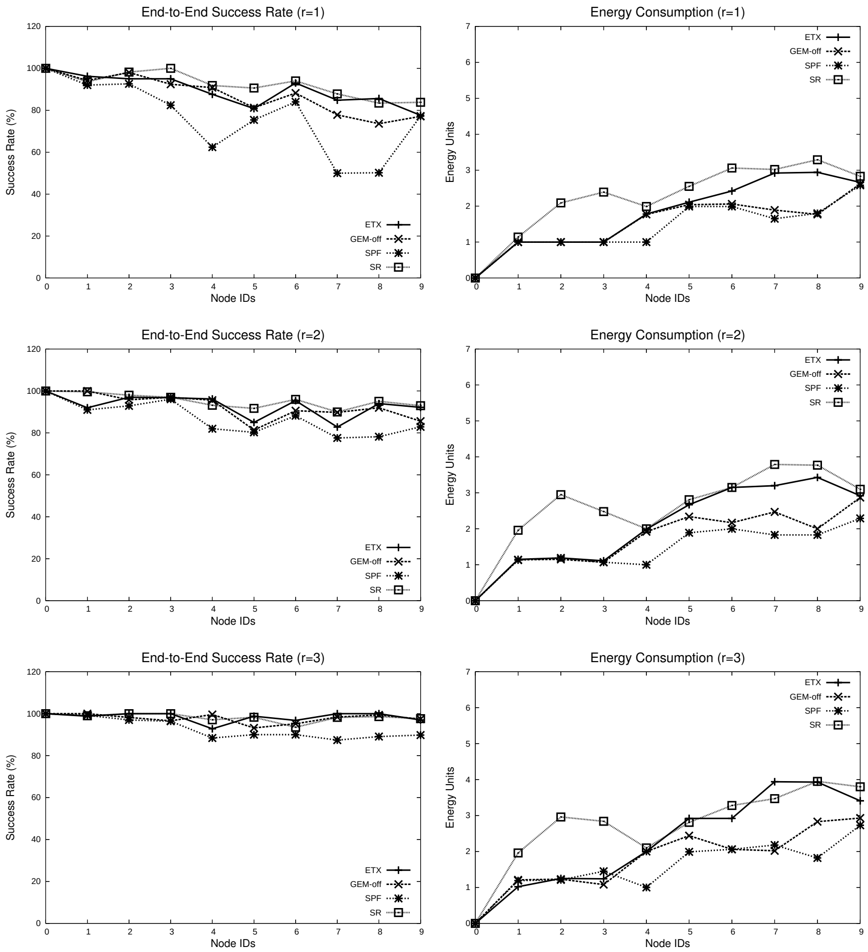

Fig. 3. End-to-end success rate and energy consumption distribution in case of 1 to 3 transmissions

However, the distribution of energy consumption shows that considering an infinite number of transmissions still leads to energy losses (according to these experiments up to 30%), because ETX considers all paths to be undirected (see formula (6)), which is not the case considering the energy demands of the path itself (see formula (4)).

Hop Distribution The graph on the left-hand side of Fig. 4 reflects the depth of the routing tree. It shows the average number of hops needed by the selected path to route packets from each node. As expected the SPF metric selects the minimal hops paths to route packets whereas the SR metric selects the longest paths. ETX and GEM behave similarly.

Route Stability The average number of parent changes during one experiment is presented in the right graph of Fig. 4 and reflects the stability of the routing tree. As can be seen, the results show a good correlation with the hop distribution metrics. The paths selected by the SPF metric are the most stable, whereas SR shows the least stability and might lead to cycles in the routing tree. ETX and GEM behave similarly and have values between SR and SPF.

Route Maintenance Overhead For all metrics, the amount of data required for route maintenance is relatively low. SPF simply forwards a hop count to the sink. Analogously, SR propagates the product of the link qualities to the sink, whereas ETX sends the expected number of transmissions. GEM requires the two values $Exp(Gain_{path})$ and $Exp(Energy_{path})$ to make its decision.

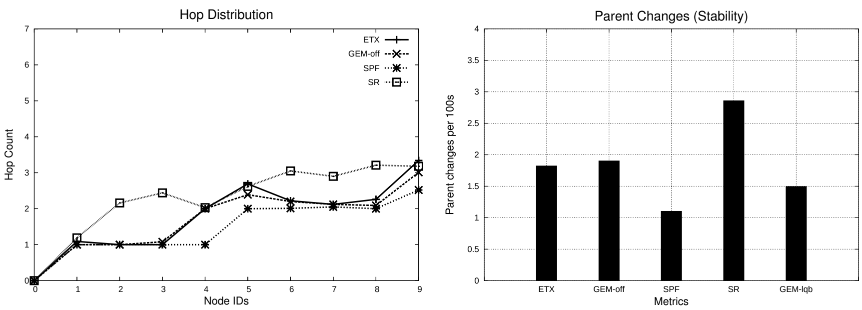

Fig. 4. Hop distribution and routing tree stability

In general, without any policy, GEM shows a slightly worse end-to-end success rate than SR and ETX but considerable gains in energy that allow it to influence the success rate if needed by the application using helper techniques like retransmissions and transmission power level tuning. Theoretically, since GEM-off tries to maximize the ratio of expected end-to-end packet success rate and energy consumption, the losses in end-to-end success rate are expected to be proportional to the gains in energy. However, in practice, link qualities are not uniformly distributed. Links from the transitional region show high variations in link quality and, therefore, the link quality estimator degrades their value as candidates to be part of the selected route in stable scenarios. This makes GEM select good quality paths even for the case where retransmissions are not available, and show comparatively high energy savings.

5.2 GEM combined with LQB

In this set of experiments we use GEM in combination with the LQB policy. The initial power level was chosen to be the same as in previous experiments $l_i = 2$. We have implemented the LQB policy to increment the transmission power level l_i by one, if at some moment p_i appears to be less than 50% and decrement l_i if the link quality is over 95%, which are some intuitive values for our scenario. We compare the GEM-lqb with GEM-off results and analyse the policy behavior in case of no retransmissions. We have used the same evaluation criteria as for the previous section, but have added convergence, which is a policy-specific criteria that indicates whether or not changes to the power level eventually stabilize.

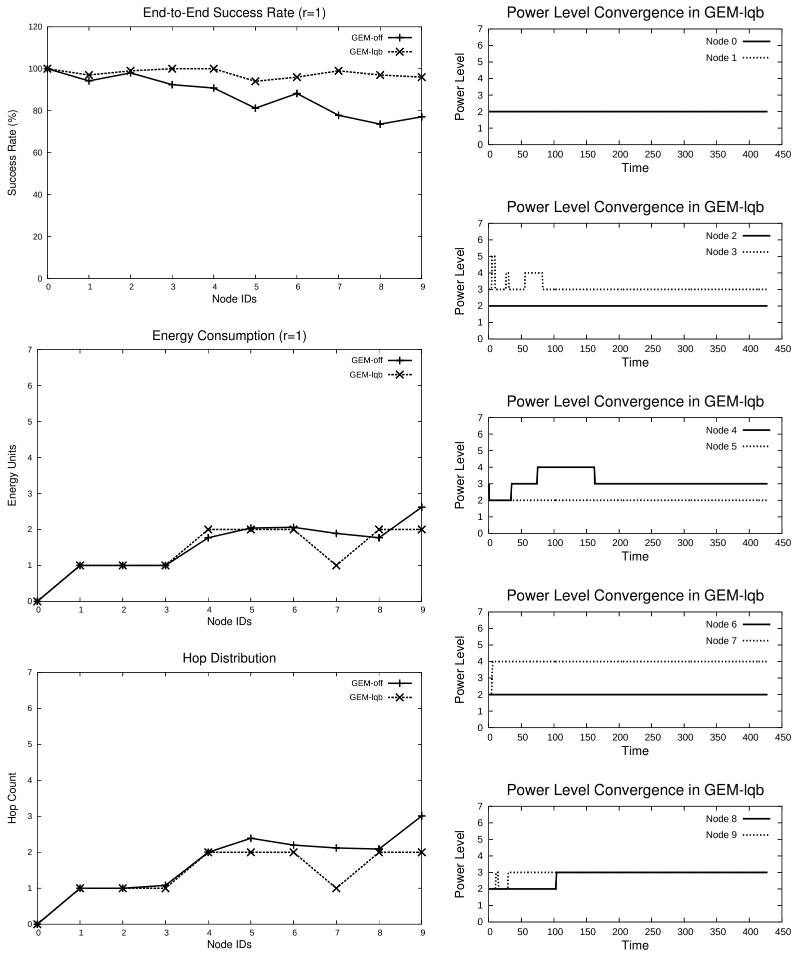

Fig. 5. Evaluation of GEM-metric with LQB policy

End-to-end Success Rate GEM-lqb shows an about 15% better success rate than GEM-off (Fig. 5, top left graph). The reason for this is the influence of the policy, which has increased the transmission power levels where needed to make individual nodes communicate with better quality.

Energy Consumption In many cases the overall evergy consumption along the path from each node to the sink is minimized even if communication now needs more energy units for one send operation (Fig. 5, left center graph). Increasing the transmission power level by one makes a send operation only up to 10% more expensive (the dependency is nonlinear on the Tmote hardware). However, it might lead to the selection of a parent node which is closer to the sink but still provides good gains. For nodes 7-9 the number of hops to the sink was minimized by increasing the transmission power range and, therefore, the energy demands were heavily reduced.

Hop Distribution Small changes in transmission power levels have minimized the depth of the routing tree to be two hops instead of three as in the previous set of experiments (Fig. 5, bottom left graph), leading to a shorter tree.

Route Stability Since the overall number of hops decreased and the selected routes were of a better quality, the average number of parent changes for GEM-lbq is lower than that of GEM-off (1.5 instead of 2, as shown in Fig. 4).

Route Maintenance Overhead The LQB policy does not generate any overhead by means of additional packets or packet fields. Its decision is solely based on local information.

Convergence In the graphs on the right-hand side of Fig. 5, we show the transmission power level of nodes $1-10$ during the first 450 seconds of the experiment. We can observe that GEM-lbq does not perform any further changes after about 160 s. However, in general convergence is not guaranteed. It might be interesting to study the conditions under which convergence can be guaranteed, but for now, we assume that it is the responsibility of the policy to ensure that this is the case.

A more detailed analysis of policy types and their specific characteristics like convergence, applicability to certain scenarios, the extents of influence on end-to-end packet success rate and energy consumption, etc. still need further research.

6 Conclusion and Future Work

In this paper we have shown that the choice of a metric for the construction of routing structures needs to explicitely take energy into account. This is especially true for long-running monitoring applications. We have proposed a new routing metric GEM and presented its mathematical foundation together with some application-specific policy examples that make it able to adapt to application needs and specifics of the environment. Our mathematical model covers most routing metrics found in the literature and present them as special cases of it. The evaluation results show that the application of GEM alone provides considerable energy savings with equivalent end-to-end packet success rate when compared to other metrics. Moreover, the resulting routing tree is naturally structured and, therefore, has good stability. GEM is parameterizable and benefits from a combination with additional application-specific policies. However, there is still a number of interesting questions that we would like to address in the future, such as a more systematic classification of policies, their convergence characteristics, connectivity, general applicability, etc. All these topics together with more thorough testing in simulation and real-life experiments are left for future work.

References

1. D. Couto, D. Aguayo, J. Bicket, R. Morris, A High-Throughput Path Metric for Multi-Hop Wireless Routing, M.I.T. Computer Science and Artificial Intelligence Laboratory, September 2003.
2. R. Draves, J. Padhye, B. Zill, Comparison of Routing Metrics for Static Multi-Hop Wireless Networks, Microsoft Research, 2004.
3. A. Woo , T. Tong , D. Culler, Taming the underlying challenges of reliable multihop routing in sensor networks, Proc. of the 1^{st} intl. conf. on Embedded networked sensor systems, November 05-07, 2003, Los Angeles, California, USA.
4. D. Son, B. Krishnamachari, J. Heidemann, Experimental study of the effects of Transmission Power Control and Blacklisting in Wireless Sensor Networks, In Proc. of SECON (2004), Santa Clara, California, USA, October 4-7, 2004.
5. O. Gnawali, M. Yarvis, J. Heidemann, R. Govindan, Interaction of Retransmission, Blacklisting, and Routing Metrics for Reliability in Sensor Network Routing, In Proc. of SECON (2004), October 2004.
6. D. D. Couto, D. Aguayo, B. Chambers, and R. Morris. Performance of multihop wireless. First Workshop on Hot Topics in Networks (HotNets-I), October 2002.
7. J. Zhao, R. Govindan, Understanding Packet Delivery Performance in Dense Wireless Sensor Networks, In Proc. of SenSys'03, Los Angeles, California, USA, 2003.
8. M. Kubisch, H. Karl, A. Wolisz, L.C. Zhong, and J. Rabaey, Distributed Algorithms for Transmission Power Control in Wireless Sensor Networks, Wireless Communications and Networking (WCNC'03), March 2003.
9. M. Carvalho, C. Margi, K. Obraczka, J. Gracia-Luna-Aceves, Modeling Energy Consumption in Single-Hop IEEE 802.11 Ad Hoc Networks, University of California Santa Cruz, 2004.
10. A. Salhieh, J. Weinmann, M. Kochhal, and L. Schwiebert, Power Efficient Topologies for Wireless Sensor Networks, ICPP 2001, September 2001.
11. C. Shurgers and M.B. Srivastava, Energy Efficient Routing in Wireless Sensor Networks, MILCOM'01, September 2001.
12. R.C. Shah and J.M. Rabaey, Energy Aware Routing for Low Energy Ad Hoc Sensor Networks, WCNC2002, March 2002.