

# Potentials of Opportunistic Routing in Energy-Constrained Wireless Sensor Networks

Gunnar Schaefer, François Ingelrest, and Martin Vetterli

Ecole Polytechnique Fédérale de Lausanne (EPFL)

Switzerland

`firstname.lastname@epfl.ch`

**Abstract.** The low quality of wireless links leads to perpetual packet losses. While an acknowledgment mechanism is generally used to cope with these losses, multiple retransmissions nevertheless occur. Opportunistic routing limits these retransmissions by taking advantage of the broadcast nature of the wireless channel: sending packets to multiple receivers at once, and only then, based on the outcome, choosing the actual next hop [1]. In this paper, we first study the potentials of opportunistic routing in energy-constrained wireless sensor networks. In particular, the reduction of retransmissions due to the broadcast advantage is balanced with the arising need for coordination to avoid duplicate packets. We then propose Coordinated Anypath Routing, an opportunistic routing protocol designed for wireless sensor networks, in which the coordination between receivers is handled by an overhearing-based acknowledgment scheme. Our protocol may be used to minimize either retransmissions or power consumption, and our simulation results show that, with lossy links, energy savings go up to 7%, even for small networks of 20 nodes.

## 1 Introduction

Multi-hop routing is a key feature of wireless ad hoc networks. Compared to traditional cellular architectures, ad hoc networks are much more flexible, as they allow users and/or devices to roam freely, without having to worry about access point locations. As long as there exists an unbroken chain of devices from a source to a destination, it is the responsibility of the routing protocol to discover it and to construct a route along it, allowing for long-distance communications.

In wireless sensor networks (WSNs), although nodes are often stationary, multi-hop routing still offers much-desired flexibility in placing, adding, and removing sensor nodes. One such example is SensorScope<sup>1</sup>, on which we have worked over the past three years. It is a time-driven WSN, used to gather dense spatio-temporal measures of various environmental parameters. Our most prominent deployment took place on top of the Généri mountain in the Swiss Alps. The gathered data allowed environmental scientists to detect and model a microclimate, which had been the cause of dangerous mud flows during strong rain

---

<sup>1</sup> <http://sensorscope.ch/>

falls [2,3]. In this particular case, multi-hop routing was a fundamental requirement, because the sink node—equipped with a GPRS transceiver—had to be placed on a nearby ridge to ensure connectivity to the GSM network. As a cellular network architecture would have forced all stations to be close to the sink, it would have prevented us from gathering the desired data.

Although much work has already been published on multi-hop routing for both ad hoc and sensor networks (*e.g.*, OLSR [4], ZRP [5], MintRoute [6]), many papers are still appearing, proposing novel algorithms, each of them trying to improve certain networking aspects (*e.g.*, energy consumption, robustness). All of these protocols, however, generally have one thing in common: at each node, a given packet is forwarded to a single, preselected neighbor. Opportunistic routing takes a different approach: each packet is forwarded to a *set* of neighbors, instead of only one [1,7]. Obviously, the probability that *at least one* node in this set receives a packet is much higher than that of a particular node receiving it. In working on the SensorScope system, we became very interested in this simple, yet intriguing idea, and have worked on adapting it to wireless sensor networks.

In this paper, we present our work on opportunistic routing in data gathering WSNs. Section 2 discusses common routing protocols for wireless ad hoc and sensor networks, while Sec. 3 provides an introduction to opportunistic routing and discusses its theoretical properties w.r.t. data gathering. Next, in Sec. 4, we show that the difficulty of efficient opportunistic routing lies in coordinating the actual receivers of a packet to avoid duplicates. For this purpose, we propose a coordinated anypath routing scheme, called CA-Path. Our simulation results show that our scheme is indeed able to minimize retransmissions when dealing with lossy links. In Sec. 5, we focus on minimizing energy consumption and show how CA-Path may be modified for this metric. Once again, our simulation results show that good energy savings may be expected with lossy links. We finally conclude in Sec. 6 and point out future work.

## 2 Routing in Wireless Ad Hoc and Sensor Networks

In ad hoc wireless networks, the limited range of radio signals forces long-distance communications to be *multi-hop*, *i.e.*, intermediate nodes between the source and the destination have to relay packets [8]. Since nodes have no *a priori* knowledge of the network topology, they have to discover it. The general idea is that nodes somehow announce their presence and monitor network traffic to learn about other nodes. After some time, each node should know about at least one route to reach each other node. Unicast, broadcast, and multicast are the typical communication primitives of ad hoc networks.

A data gathering WSN is a special case of an ad hoc network, typically comprising a large number of data sources, but only a small number of data sinks, or even just a single one [9]. The responsibilities of the routing protocol are thus limited to ensuring data-flow to a very small number of destinations, as opposed to all other nodes in the network. Most often, data-gathering networks are abstracted as sink-rooted trees.

A large number of routing protocols exists for both ad hoc and sensor networks. These protocols are commonly classified according to various criteria, for instance, whether route discovery is reactive (*e.g.*, DSR [10]), or proactive (*e.g.*, OLSR [4]). Orthogonally, protocols may be singlepath, such as MintRoute [6] and the Collection Tree Protocol [11], or multipath (*e.g.*, Maximum Lifetime Routing [12]). Many more classification criteria may, in fact, be defined [13].

Nevertheless, all of these criteria fail to highlight the fact that forwarding packets is always performed in the same way: the sender selects a single next hop according to a given metric, and packets are sent exclusively to that node. This behavior actually stems from the well-established wired networks and, as we will elaborate in the next section, may not be ideally suited for *wireless* networks.

### 3 Opportunistic Routing for Data Gathering


As mentioned in the introduction, opportunistic routing views the routing problem with a different philosophy, compared to traditional algorithms. While choosing a next hop *a priori*, *i.e.*, before actually sending a packet, may be well suited for wired networks (in which losses are almost inexistent), this is not necessarily the case for wireless networks. Due to the broadcast nature of the wireless channel, packets are, in fact, typically transmitted to all nearby nodes<sup>2</sup>, regardless of any chosen next-hop. Hence, finalizing the next hop routing decision only *after* sending a packet to a set of possible relays may be a better solution.

For instance, in Fig. 1, when Node *E* has to send a packet to the *Sink*, its three possible next hops are *A*, *B*, and *C*. By preselecting one of them, the probability that the chosen node will not receive the packet—and thus the retransmission probability—is  $1 - 0.4 = 0.6$ . Node *E*, however, should not have to care about which node will be the actual next hop. Much rather, it could simply broadcast the packet to the *candidate relays* *A*, *B*, and *C*. Subsequently, an agreement will have to be reached, regarding which node(s) will forward the packet. In this case, the probability that at least one of the prospective relays receives the packet is  $1 - (1 - 0.4)^3 = 0.784$ . The retransmission probability of the packet is now equal to only 0.216, much lower than the original 0.6. Moreover, even if one of the links was very lossy, opportunistic routing could make use of it, resulting in a natural load balancing between *A*, *B*, and *C*. Most traditional routing algorithms would use only the better links in such a scenario.

Generalizing the Bellman-Ford algorithm, the cost of an opportunistic route comprises two components: the *anycast link cost* (ALC), which is the cost to reach the *candidate relay set* (CRS), and the *remaining path cost* (RPC), which is the cost to get from the candidate relay set to the destination. Following our previous example, the ALC is the cost to send a packet from *E* to the set  $\{A, B, C\}$ , and the RPC is the cost to send the packet from this set to the *Sink*.

---

<sup>2</sup> We assume customary omnidirectional antennas.



**Fig. 1.** Topology of a wireless network. All links have a delivery probability of 0.4.

As it is our aim to reduce the number of retransmissions, we draw upon the expected transmission count metric (ETX) [14]. In the following, the term “cost” will thus always imply the ETX metric.

### 3.1 Anypath Routing

Anypath routing provides a general way of computing the cost to reach a given destination [7]. To find the *shortest anypath route* (SAR) between two nodes, we assign a CRS to each node, such that the expected cost of forwarding packets via this set to the destination is minimized. Anypath routing provably computes the optimal set for each node to reach a certain destination. For this, Dubois-Ferrière *et al.* have proposed the following three equations [7].

Node  $i$ 's anycast link cost, which accounts for reaching at least one node in the CRS  $J$ , is defined as

$$\text{ALC}_{iJ} = \frac{1}{p_{iJ}} = \frac{1}{1 - \prod_{j \in J} (1 - p_{ij})}, \quad (1)$$

where  $p_{ij}$  is the probability of successful transmission between nodes  $i$  and  $j$ , and  $p_{iJ}$  is the probability of successful transmission between  $i$  and at least one node in the set  $J$ .

The remaining path cost of node  $i$ 's CRS  $J$  is defined as

$$\text{RPC}_{iJ} = \frac{C_1 p_{i1} + \sum_{j=2}^{|J|} C_j p_{ij} \prod_{k=1}^{j-1} (1 - p_{ik})}{1 - \prod_{j \in J} (1 - p_{ij})}, \quad (2)$$

assuming that the nodes in  $J$  are sorted by their cost, *i.e.*,  $C_1 < C_2 < \dots < C_{|J|}$ . The numerator of this equation represents the probability of a packet being received by a particular node in the CRS  $J$  and not being received by any node with a lower cost to reach the destination, while the denominator accounts for the probability that at least one node in  $J$  has received the packet.

The cost of the shortest anypath route is defined as

$$C_i^{\text{SAR}} = \min_{J \in 2^{\mathcal{N}(i)}} (\text{ALC}_{iJ} + \text{RPC}_{iJ}), \quad (3)$$

where  $N(i)$  is the set of neighbors of node  $i$ , and  $2^{N(i)}$  is the power set of  $N(i)$ . Due to the aforementioned sorting of  $J$ , the search space is only of size  $n = |N(i)|$ , as opposed to the  $2^{|N(i)|} - 1$  non-empty subsets of  $i$ 's neighbors. If it is not beneficial to add neighbor  $j$  to the CRS, then it cannot be beneficial to add neighbor  $j + 1$ , with cost greater than that of  $j$ . Furthermore, the cost of the shortest anypath route will never be higher than that of the shortest singlepath route, since the set of opportunistic routes between two nodes includes all singlepath routes.

In examining these equations, an inherent tension between ALC and RPC becomes apparent. The ALC is minimized by taking the entire neighborhood as candidate relay set; on the contrary, the RPC is minimized by choosing only the one neighbor with the least cost to reach the destination. Note also that these equations do not consider the acknowledgment cost, although acknowledging packets is mandatory to prevent losses.

### 3.2 Anypath Routing for Data Gathering

Anypath routing has been specifically targeted towards ad hoc networks, in which any pair of nodes may wish to communicate. This stands in contrast to data gathering WSNs, in which communication occurs only between the sensor nodes and the sink. However, as long as there is no in-network data processing (*e.g.*, aggregation), data gathering represents a special case, in which computational overhead and memory footprint are both reduced, as each node must maintain the optimal CRS for only a single destination.

Returning to the example network of Fig. 1, we now compare the cost of the shortest singlepath route to that of the shortest anypath route for the particular case of Node E.

For both singlepath and anypath, the cost to send a packet from  $A$ ,  $B$ , or  $C$  to the *Sink* is

$$C_A = C_B = C_C = \frac{1}{0.4} = 2.5.$$

*Singlepath:* The cost to send a packet from  $E$ , via  $A$  (or via either  $B$  or  $C$ ), to the *Sink* is

$$C_E = \frac{1}{0.4} + C_A = 5.0.$$

*Anypath:* The cost to send a packet from  $E$ , via the CRS  $J(E) = \{A, B, C\}$ , to the *Sink* is

$$\begin{aligned} C_E &= \text{ALC}_{E,\{A,B,C\}} + \text{RPC}_{E,\{A,B,C\}} \\ &= \frac{1}{1 - (1 - 0.4)^3} + \frac{2.5 \cdot 0.4 + 2.5 \cdot 0.4 \cdot (1 - 0.4) + 2.5 \cdot 0.4 \cdot (1 - 0.4)^2}{1 - (1 - 0.4)^3} = 3.78. \end{aligned}$$

Table 1, providing the results for the entire network, clearly shows the potential of anypath routing. Note that the given network is an intentionally simple

**Table 1.** Comparison of singlepath and anypath routing w.r.t. Fig. 1

<u>Node</u>	<u>Singlepath</u>	<u>Anypath</u>	<u>Improvement</u>
<i>A</i>	2.50	2.50	0.0%
<i>B</i>	2.50	2.50	0.0%
<i>C</i>	2.50	2.50	0.0%
<i>D</i>	5.00	4.01	19.8%
<i>E</i>	5.00	3.78	24.5%
<i>F</i>	5.00	4.01	19.8%
<i>G</i>	7.50	5.17	31.1%
<b>Total</b>	30.00	24.46	18.5%

example; the advantage of anypath becomes more pronounced as the number of hops increases. Note also that we focus on shortest singlepath routing for comparison. Other protocols would show similar results, as all of them use a single next hop and would thus face the same losses.


## 4 Coordinated Anypath Routing

So far, we have implicitly assumed the existence of a mechanism ensuring that only one receiver—the one with the lowest cost—will indeed forward a packet. Without such a mechanism, many duplicate packets will appear, resulting in substantial growth of the energy consumption. For instance, considering Fig. 1, the full cost of sending a packet from Node *E*, via the CRS  $\{A, B, C\}$ , to the *Sink* would be  $C_E = 5.1$ , even more than the cost of the shortest singlepath route. This cost accounts for all possible receiver scenarios: (i) one candidate receives, (ii) two candidates receive, and (iii) all three candidates receive (and forward) the packet. Hence, coordinating the actual receivers is mandatory.

### 4.1 Receiver Coordination

A few solutions for receiver coordination have been proposed. While cooperative diversity schemes [15] are information-theoretically interesting, they are incompatible with today's WSN hardware. RTS/CTS-based methods (*e.g.*, MAC-layer anycasting [16]) require additional messages to be sent, that can quickly outweigh the gains obtained thanks to opportunistic routing.

Extremely Opportunistic Routing (ExOR), designed for throughput maximization [1], comprises an overhearing-based coordination scheme. To choose the effective next hop, the sender includes in its packets a prioritized list of the CRS members. Next, the receivers send their acknowledgments (ACKs) in a staggered fashion (see Fig. 2), based on each node's position in the aforementioned list. As the nodes listen to each other, they include, in their own ACK, the ID of the highest-priority actual receiver they know about—possibly, their own ID. Then, all nodes believing to be the highest-priority receiver further relay the packet. The original sender of the packet considers it successfully forwarded



**Fig. 2.** The ordered list of intended receivers, sent as part of each packet, fosters collisions avoidance: ACKs are sent in a staggered fashion, rather than concurrently. Nodes include the ID of the highest-priority actual receiver they know about (by overhearing) in their own ACK. This example assumes that  $A$  and  $C$  cannot hear each other, but  $B$  can communicate with  $A$  and with  $C$ . Thus, both  $B$  and  $C$  learn about a higher-priority receiver, and drop their packets accordingly, while  $A$  will forward its packet.

as soon as it receives one ACK. Obviously, the emergence of multiple forwarders is not entirely eliminated, as it is not guaranteed that receivers are sufficiently able to overhear each other.

ExOR has been designed to increase throughput, not for energy-efficient data gathering, and we cannot use it directly for our purpose. However, as it is our goal to reduce retransmissions, we have chosen to follow the same idea for coordinating receivers. Not only does it avoid additional transmissions (ACKs are anyhow mandatory to prevent losses), but it relies solely on adding a few bytes to packets, which has little impact on transmissions (see, for instance, Fig. 3 in Sec. 5, showing the power consumption of a typical sensor mote).

## 4.2 Coordinated Anypath Routing for Data Gathering

Based on the cost of the shortest anypath route, given in Equation (3), the cost of the shortest *coordinated* anypath route (SCA-Path) is equal to

$$C_i^{\text{SCA-Path}} = \min_{J \in 2^{N(i)}} (\text{ALC}_{i,J} + \text{AC}_{i,J} + \text{RCC}_J + \text{RPC}_{i,J}), \quad (4)$$

where  $\text{AC}_{i,J}$  is the acknowledgment cost of CRS  $J$  at node  $i$ , and  $\text{RCC}_J$  is the cost of coordination among the nodes in CRS  $J$ .

As we are using a purely overhearing-based coordination approach, the  $\text{RCC}$  can immediately be set to zero:

$$C_i^{\text{SCA-Path}} = \min_{J \in 2^{N(i)}} (\text{ALC}_{i,J} + \text{AC}_{i,J} + \text{RPC}_{i,J}), \quad (5)$$

where the respective definitions of  $\text{ALC}_{i,J}$  and  $\text{RPC}_{i,J}$  remain as given in Sec. 3.

To define  $\text{AC}_{i,J}$ , we must introduce some additional notation. Let  $\mathbf{A} \in \{0, 1\}^{|J(i)|}$  be a random vector representing the outcome of the current transmission. Let  $\mathbf{A}_j = 1$  iff  $j \in A(i)$ , where  $A(i)$  is the set of actual receivers; hence,  $A(i) \subseteq J(i)$ . Finally,  $a : \{0, 1\}^{|J(i)|} \rightarrow \mathbb{R}$  is the function, which assigns—to a specific realization of  $\mathbf{A}$ —the acknowledgment cost, with  $a(\mathbf{0}) = 0$ .  $\text{AC}_{i,J}$  is thus the expected cost of acknowledging a packet by a certain set to actual receivers, conditional on at least one node in the CRS receiving the packet:

**Table 2.** Comparison of singlepath and CA-Path routing w.r.t. Fig. 1

<u>Node</u>	<u>Singlepath</u>	<u>CA-Path</u>	<u>Improvement</u>
<i>A</i>	5.00	5.00	0.0%
<i>B</i>	5.00	5.00	0.0%
<i>C</i>	5.00	5.00	0.0%
<i>D</i>	10.00	9.69	3.1%
<i>E</i>	10.00	9.69	3.1%
<i>F</i>	10.00	9.69	3.1%
<i>G</i>	15.00	14.38	4.2%
<b>Total</b>	60.00	58.44	2.6%

$$AC_{i,J} = E \left[ a(\mathbf{A}) \mid \sum_{k \in J} \mathbf{A}_k > 0 \right] = \frac{E[a(\mathbf{A})]}{P\left(\sum_{j \in J} \mathbf{A}_j > 0\right)} = \frac{\sum_{A \in 2^J} P(A) a(A)}{1 - \prod_{j \in J} (1 - p_{ij})}, \quad (6)$$

where  $2^J$  is the power set of  $J$ , and  $P(A)$  is the probability that the set of actual receivers is indeed  $A$ :

$$P(A) = \prod_{j \in A} p_{ij} \cdot \prod_{j \in J \setminus A} (1 - p_{ij}), \quad (7)$$

and  $a(A)$  is the cost of acknowledging a packet received by all nodes in  $A$ :

$$a(A) = \sum_{j \in A} \frac{1}{p_{ji}}. \quad (8)$$

We call this coordinated anypath routing scheme *CA-Path*. From hereon, singlepath costs will include acknowledgment costs, just as CA-Path costs do. Note that the optimal CRS determined by CA-Path is likely to be smaller than that found with anypath routing, as a larger CRS incurs a higher acknowledgment cost. Furthermore, similar to anypath routing, the cost of the shortest CA-Path route will never be higher than that of the shortest singlepath route.

Returning to Fig. 1, we now compare the cost of the shortest singlepath route to that of the shortest CA-Path route.

For both schemes, the cost to send a packet from  $A$ ,  $B$ , or  $C$  to the *Sink* is

$$C_A = C_B = C_C = \frac{1}{0.4} + \frac{1}{0.4} = 5.0.$$

*Singlepath:* The cost to send a packet from  $E$ , via  $A$  (or via either  $B$  or  $C$ ), to the *Sink* is

$$C_E = \frac{1}{0.4} + \frac{1}{0.4} + C_A = 10.0.$$



*CA-Path*: The cost to send a packet from  $E$ , via the CRS  $J(E) = \{A, B\}$ , to the *Sink* is

$$\begin{aligned} C_E &= \text{ALC}_{E,\{A,B\}} + \text{AC}_{E,\{A,B\}} + \text{RPC}_{E,\{A,B\}} \\ &= \frac{1}{1 - (1 - 0.4)^2} + \frac{2 \cdot \frac{1}{0.4} \cdot 0.4 \cdot (1 - 0.4) + \left(\frac{1}{0.4} + \frac{1}{0.4}\right) \cdot 0.4^2}{1 - (1 - 0.4)^2} \\ &\quad + \frac{5.0 \cdot 0.4 + 5.0 \cdot 0.4 \cdot (1 - 0.4)}{1 - (1 - 0.4)^2} = 9.69. \end{aligned}$$

This is the cost of the shortest CA-Path route from Node  $E$  to the *Sink*. Note that this optimal CRS is smaller than the one resulting from the original anypath equations. Table 2 provides the results for all the nodes of Fig. 1. Although we are now considering acknowledgment costs, CA-Path still has an advantage.

### 4.3 Theoretical Bounds

Let us assume a network, in which all links have the same delivery probability  $p$ . For such a network, we can determine the threshold probability, below which selecting multiple next hops with CA-Path provides better results than singlepath. Let us consider Node  $D$  of the network shown in Fig. 1: its two possible next hops towards the *Sink* are  $A$  and  $B$ . Since all probabilities are equal,  $D$  obviously cannot consider  $E$  and  $G$  as candidate next hops. The question, of whether it is better to use two next hops instead of one, reduces to solving the following inequality:

$$C_{D,\{A,B\}}^{\text{CA-Path}} < C_{D,\{A\}}^{\text{SP}}. \quad (9)$$

In other words, for which values of  $p$  can CA-Path decrease the cost compared to singlepath? Developing this equation, we find

$$\frac{1}{1 - (1 - p)^2} + \frac{2 \cdot \frac{1}{p} \cdot p \cdot (1 - p) + \frac{2}{p} \cdot p^2}{1 - (1 - p)^2} + \frac{\frac{2}{p} \cdot p + \frac{2}{p} \cdot p \cdot (1 - p)}{1 - (1 - p)^2} < \frac{2}{p} + \frac{2}{p},$$

which solves to  $p < 0.5$ . Thus, when  $p \geq 0.5$ , CA-Path reduces to singlepath. Following similar reasoning, choosing three or even four next hops is interesting only when  $p < 0.38$  or  $p < 0.31$ , respectively. This explains why, in Fig. 1, where  $p = 0.4$ , Node  $E$  uses only two candidate relays.

### 4.4 Simulation Results

In order to concentrate on the networking aspects, and avoid issues unrelated to CA-Path, we have developed our own open-source simulation tool<sup>3</sup>, rather than relying on one of the existing heavyweight frameworks. The ns-2 simulator<sup>4</sup>, for

<sup>3</sup> <http://rr.epfl.ch/19/>

<sup>4</sup> <http://nslam.isi.edu/nsnam/>

**Table 3.** ETX performance of CA-Path for different types of networks:  $n$  is the number of nodes and  $d$  the average maximum number of hops to the sink

<u>Network</u>	<u>Singlepath</u>	<u>CA-Path</u>	<u>Improvement</u>
$n = 20, d = 2.6$	250.80	241.68	3.6%
$n = 35, d = 3.6$	566.67	540.76	4.6%
$n = 50, d = 4.4$	956.75	906.80	5.2%
$n = 100, d = 6.1$	2663.33	2507.12	5.9%
$n = 250, d = 9.4$	10152.46	9483.79	6.6%
$n = 500, d = 12.9$	27915.92	25982.50	6.9%

instance, entails great effort for an in-depth understanding of the interactions between its numerous components. Moreover, it requires so many modifications and add-ons for simple primitives (*e.g.*, broadcasting a packet to all neighbors) that results may not be trustworthy in the end.

Each result provided in this section is an average over 500 generated network topologies. For each topology, nodes are uniformly and randomly placed in a square area, whose size is determined by the number of nodes and the average node degree, which is fixed at 10. A single sink is present in each network. Non-connected topologies are discarded and regenerated as necessary. Link quality estimation is out of the scope of this paper; we assume that a neighborhood discovery protocol is in charge of it. A simple solution, such as counting sequence numbers, is sufficient for this purpose [2].

To limit computational complexity, the maximum CRS size in is set to three, *i.e.*, nodes select at most three next hops, even if more would further decrease their cost. If we assume that node identifiers are stored on a single byte, this leads to a maximum overhead of three bytes per data packet, which is negligible w.r.t. energy consumption. More on this subject is elaborated in the next section.

As we have pointed out previously, using CA-Path is obviously sensible only when links are lossy. In this case, the number of additional ACK transmissions is less than the number of data packet retransmissions and thus leads to a smaller ETX. On the contrary, when links are strong, CA-Path reduces to singlepath. Hence, for the results presented in this section, all links in the generated networks have a delivery probability of  $p = 0.25$ .

Table 3 provides the average overall ETX for various network sizes. We have considered many scenarios: small, medium, large, and very large networks. Most current WSN deployments are rather small-scale (*e.g.*, LUSTER [17], SensorScope [2]), but larger scenarios are envisioned for the near future. As expected, savings increase with the average distance to the sink, since CA-Path is able to save a few transmissions over singlepath at each hop. In fact, the last hop to the sink limits the savings, since it is cheaper for many nodes to communicate exclusively with the sink. In this case, CA-Path resumes to singlepath. Overall, results are promising for all kinds of networks, as savings range from 3.6% for small networks of 20 nodes to 6.9% for very large networks of 500 nodes. This clearly shows the potential of CA-Path when dealing with lossy links.

## 5 Minimizing Energy Consumption

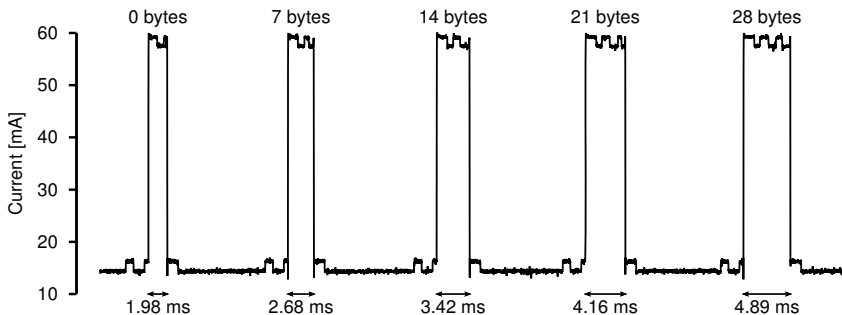
So far, we have only considered the problem of minimizing the number of transmissions, while—in the real world—minimizing energy consumption is generally more interesting. These two metrics are of course correlated, but even with a fixed transmission power, sending packets of different sizes results in different energy consumption.

To determine the respective costs of ACKs and data packets, we have measured the power consumption of a TinyNode sensor mote<sup>5</sup>, while transmitting packets of different payload lengths (see Fig. 3). The power consumption when sending a 0-byte payload is not zero, because each payload is preceded by a network header. Moreover, the radio precedes each packet with a specific pattern of bits, so that receivers can detect the beginning of incoming packets<sup>6</sup>. Based on our measurements, the duration of transmission—and thus the power consumption—may be approximated as

$$C(l) = 0.1043 \cdot l + 1.966, \quad (10)$$

where  $l$  is the payload length, in bytes. According to this equation, sending a 28-byte data packet costs  $C(28) = 4.89$  while a 2-byte acknowledgment costs  $C(2) = 2.17$ . The ratio between the two costs is 0.44, which we approximate by 0.5, making ACKs a bit more costly than in reality. Thus, in the following, we assume that an acknowledgment consumes half as much energy as a data packet.

Considering these relative costs has no impact on singlepath, which will always select the same path to send packets to the sink. With CA-Path, things are different, as the CRS size impacts the acknowledgment cost. To take this into account, we must modify Equation (8), such that  $a(A)$  is now equal to



**Fig. 3.** Power consumption of a TinyNode mote when sending packets of various payload lengths. The transmission power is set to 15 dBm.

<sup>5</sup> <http://tinynode.com/>

<sup>6</sup> Note that this is not related to a low power listening mechanism.

$$a(A) = \sum_{j \in A} \frac{0.5}{p_{ji}}, \quad (11)$$

while the other cost functions remain unchanged.

## 5.1 Implementation Issues

CA-Path may be implemented as a purely overhearing-based, proactive routing protocol, *i.e.*, no additional messages are required for route maintenance. Nodes must keep an up-to-date table of their neighbors' costs and include their own cost to reach the sink, together with the chosen list of intended receivers, in each data packet. Any node, overhearing this information, can update its neighborhood table accordingly. To get the process started, the sink must send out beacons, advertising its own cost of zero.

In order to find the best CRS, each node needs to examine all possible subsets of its neighbors. Let us assume that  $d$  is the number of neighbors and  $m$  is the maximum CRS size ( $1 \leq m \leq d$ ) we wish to consider. With the full-fledged scheme presented above, the number of candidate relay sets to examine is

$$\sum_{k=1}^m \binom{d}{k},$$

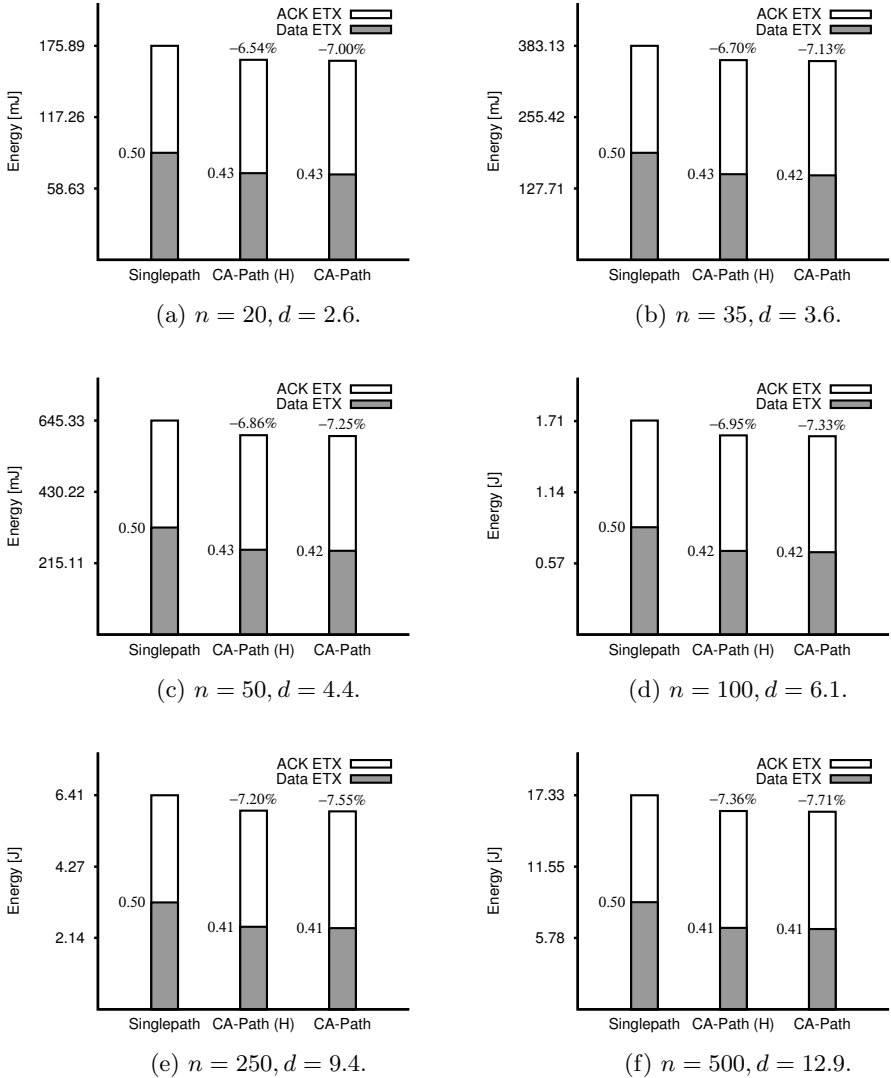
which results in a time complexity of  $O(d^m)$ . As savings with CA-Path may occur only when  $m \geq 2$ , complexity quickly becomes a problem on today's nodes.

To overcome this issue, we propose the following greedy heuristic: each node evaluates its own cost w.r.t. all singleton CRSs and sorts those accordingly. Now, the node sets its CRS to be the least expensive singleton and tries to merge it with the next best one. If that decreases its cost, the node sets its CRS to be these two nodes and then tries to merge it with the third best singleton. The process is repeated until (i) the cost of the current node no longer decreases or (ii) the CRS has reached its maximum size.

This heuristic will consider at most  $d+m-1$  CRSs, leading to a complexity of  $O(d)$ , which is much better suited to the capabilities of current sensor nodes, as well as being more scalable. As we will experimentally show below, the heuristic is a good approximation of the exhaustive search for the optimal CRS.

## 5.2 Simulation Results

To evaluate our protocol in terms of energy consumption, we have used the same simulation tool as before, incorporating the modified equation for CA-Path. All parameters remain the same, except for the generated topologies: instead of using a uniform delivery probability, each link is assigned a random probability  $p$ , such that  $0.1 \leq p \leq 0.4$ . This modification is needed to evaluate the heuristic (denoted CA-Path (H)) we have proposed above; with uniform probabilities, there is no difference to the full-fledged scheme.



**Fig. 4.** Performance of CA-Path when minimizing energy consumption. The “(H)” indicates use of the heuristic instead of exhaustive search for the optimal CRS. Each bar also shows the ratio between data packet and ACK transmissions.

Figure 4 depicts the average overall expected energy consumption for the various network sizes we have considered. This consumption includes the transmission (and acknowledgment) of one data packet from each node to the sink. For instance, when  $n = 20$ , using singlepath leads to a global consumption of 175.89mJ at each transmission cycle. To compute these values, we made use of the TinyNode power consumption data illustrated in Fig. 3. Sending a

28-byte data packet draws 60 mA at 3.3 V for 4.89 ms, leading to a consumption of 0.97 mJ. Similarly, sending an ACK consumes 0.43 mJ.

From these results, we can see that the ratio of energy saved by CA-Path ranges from 7.00% for small networks of 20 nodes to 7.71% for large networks of 500 nodes. Overall, the greedy heuristic provides results close to the optimal ones, always within half of a percentage point. This is encouraging, since the complexity of the heuristic is linear, making it very scalable.

Figure 4 also provides the distribution of the expected number of transmissions between data packets and acknowledgments. For singlepath, the ratio is always 0.5, since each data packet triggers one ACK. For CA-Path, we can observe that the transmission load is shifted from data packets to ACKs, and the ratio goes down to 0.41 for the 500-node networks. As a result, even when the same overall number of packets is sent, CA-Path leads to energy savings due to the smaller cost of ACKs.

## 6 Conclusion and Perspectives

In this paper, we have studied the potentials of opportunistic routing in data gathering wireless sensor networks. CA-Path, our implementation of a coordinated opportunistic routing scheme, is specifically designed to limit energy wasting when dealing with radio links of poor quality. We have shown that with strong links, CA-Path reduces to singlepath, while, when working in difficult environments with lossy links (*e.g.*, very long distance between nodes), CA-Path is a viable solution, leading to energy savings. With its linear time complexity, the heuristic we have proposed is especially suited to embedded sensor motes and provides results very close to exhaustive search.

The results we have shown must, however, be moderated a bit: in our simulations, routes are loop-free, and all actual receivers of a data packet can hear each other. While loops may affect any routing protocol and measures may be taken against them, non-overheard ACKs are more difficult to cope with. Due to our selection scheme, next hops should be close to each other, but of course this does not ensure that they can actually hear each other. The resulting duplicates have the potential to outweigh the savings obtained with CA-Path. We are thus planning to benchmark CA-Path by implementing it on a real sensor network. We will also study the extend of load balancing induced by CA-Path, in comparison to traditional protocols. In this area, too, good results are expected.

Finally, our results pave the way to even greater energy savings, when combined with other coordination schemes. Decreasing the coordination cost will lead CA-Path to using larger candidate relays sets and thus to lower energy consumption. One possibility could be to acknowledge multiple data packets at once, by working with a window of sequence numbers, similar to TCP. In this case, the cost of ACKs would be lowered, although at the expense of higher latency. Hence, application-specific solutions should provide the best results, overall.

## Acknowledgments

This work was partially financed by the Swiss National Center of Competence in Research for Mobile Information and Communication Systems (NCCR MICS) and the European Commission under the FP6 project WASP.

## References

1. Biswas, S., Morris, R.: ExOR: Opportunistic multi-hop routing for wireless networks. In: Proceedings of the ACM SIGCOMM Conference (October 2005)
2. Barrenetxea, G., Ingelrest, F., Schaefer, G., Vetterli, M., Couach, O., Parlange, M.: Sensorscope: Out-of-the-box environmental monitoring. In: Proceedings of the ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN) (April 2008)
3. Barrenetxea, G., Ingelrest, F., Schaefer, G., Vetterli, M.: The hitchhiker's guide to successful wireless sensor network deployments. In: Proceedings of the ACM International Conference on Embedded Networked Sensor Systems (SenSys) (November 2008)
4. Jacquet, P., Mühlethaler, P., Clausen, T., Laouiti, A., Qayyum, A., Viennot, L.: Optimized link state routing protocol for ad hoc networks. In: Proceedings of the IEEE International Multi-topic Conference (INMIC) (December 2001)
5. Haas, Z.J.: A new routing protocol for the reconfigurable wireless networks. In: Proceedings of the IEEE International Conference on Universal Personal Communications (October 1997)
6. Woo, A., Tong, T., Culler, D.: Taming the underlying challenges of reliable multi-hop routing in sensor networks. In: Proceedings of the ACM International Conference on Embedded Networked Sensor Systems (SenSys) (November 2003)
7. Dubois-Ferrière, H., Grossglauser, M., Vetterli, M.: Least-cost opportunistic routing. In: Proceedings of the Allerton Conference on Communication, Control, and Computing (September 2007)
8. Basagni, S., Giordano, S., Stojmenović, I.: Mobile Ad Hoc Networking. IEEE Computer Society Press, Los Alamitos (2004)
9. Estrin, D., Girod, L., Pottie, G., Srivastava, M.: Instrumenting the world with wireless sensor networks. In: Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP) (May 2001)
10. Johnson, D.B.: Routing in ad hoc networks of mobile hosts. In: Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications (WMCSA) (December 1994)
11. Fonseca, R., Gnawali, O., Jamieson, K., Kim, S., Levis, P., Woo, A.: The collection tree protocol (CTP) (2006), <http://www.tinyos.net/tinyos-2.x/doc/html/tep123.html>
12. Chang, J.H., Tassiulas, L.: Maximum lifetime routing in wireless sensor networks. IEEE/ACM Transactions on Networking 12(4), 609–619 (2004)
13. Al-Karaki, J.N., Kamal, A.E.: Routing techniques in wireless sensor networks: A survey. IEEE Wireless Communications 11(6), 6–28 (2004)
14. Couto, D.D., Aguayo, D., Bicket, J., Morris, R.: A high-throughput path metric for multi-hop wireless routing. In: Proceedings of the ACM International Conference on Mobile Computing and Networking (MobiCom) (September 2003)

15. Laneman, J., Tse, D., Wornell, G.: Cooperative diversity in wireless networks: Efficient protocols and outage behavior. *IEEE Transactions on Information Theory* 50(12), 3062–3080 (2004)
16. Choudhury, R., Vaidya, N.: Mac-layer anycasting in ad hoc networks. In: *Proceedings of the ACM SIGCOMM Conference* (August 2004)
17. Selavo, L., Wood, A., Cao, Q., Sookoor, T., Liu, H., Srinivasan, A., Wu, Y., Kang, W., Stankovic, J., Young, D., Porter, J.: LUSTER: Wireless sensor network for environmental research. In: *Proceedings of the ACM International Conference on Embedded Networked Sensor Systems (SenSys)* (November 2007)