SOFA: Communication in Extreme Wireless Sensor Networks

Marco Cattani*, Marco Zuniga, Matthias Woehrle, and Koen Langendoen

Delft University of Technology, Delft, The Netherlands,
{m.cattani, m.zuniga, matthias.woehrle, k.g.langendoen}@tudelft.nl

Abstract. Sensor networks can nowadays deliver 99.9% of their data with duty cycles below 1%. This remarkable performance is, however, dependent on some important underlying assumptions: low traffic rates, medium size densities and static nodes. In this paper, we investigate the performance of these same resource-constrained devices, but under scenarios that present extreme conditions: high traffic rates, high densities and mobility. To cope with these stringent requirements, we propose a novel communication protocol named SOFA (Stop On First Ack). SOFA utilizes opportunistic anycast to drastically reduce the rendezvous times of asynchronous duty cycled nodes –long rendezvous times are the key limitation of protocols operating under high densities and high traffic conditions. SOFA is also stateless, which makes it resilient to mobility. We implemented SOFA in the Contiki OS and tested it both in simulation and on a 100-node testbed. Our results show that SOFA reliably communicates in mobile networks with extreme densities (hundreds of nodes) and higher traffic rates (packets per second) while maintaining a low duty cycle (≈2%). Under these extreme conditions, current duty cycled protocols collapse.

1 Introduction

Hitherto, the protocol stack of wireless sensor networks has been mainly designed and optimized for applications satisfying one or more of the following conditions: (i) low traffic rates (a packet per node every few minutes), (ii) medium sized densities (tens of neighbors), and (iii) static topologies. There are, however, many scenarios where these relatively mild network conditions do not hold and traditional protocol stacks simply collapse. We consider mobile networks consisting of hundreds of thousands of nodes with hundreds of neighbors that need to disseminate information at a relatively high rate (a packet per node every few seconds, instead of every few minutes).

Our work is motivated by a scenario related to public safety: the need to monitor crowds in open-air festivals. SOFA is part of the EWiDS project aimed at providing attendees with coin-size devices that can actively monitor the density of their surroundings and issue alerts when crossing dangerous thresholds.

* Marco Cattani was supported by the Dutch national program COMMIT\

From a networking perspective, communication in mobile scenarios with high traffic rates and high densities poses several non-trivial technical challenges. First, similar to traditional WSN, these Extreme Wireless Sensor Networks (EWSN) also work with devices with limited energy resources and need to rely on radio duty-cycling techniques to save energy. Second, due to the network scale and node mobility, we cannot rely on methods that combine duty cycling techniques with central coordinators [8] or that require some level of synchronization between the wake up periods of a given node and its neighbors [7, 17]. The system must be asynchronous and fully distributed. Third, due to their inefficient bandwidth utilization, traditional unicast and broadcast primitives –which are asynchronous, distributed and built on top of duty cycling techniques [3, 16]– simply collapse under the traffic-demands of EWSN.

Henceforth, providing an energy-efficient communication primitive in EWSN requires a careful evaluation of the following problem: in asynchronous duty cycling techniques, much of the bandwidth is wasted in coordinating the rendezvous of the (sleeping) nodes. In EWSN, nodes need to reduce this overhead to free up the channel's bandwidth for actual *data* transmissions.

To tackle this problem, SOFA (*Stop on First Ack*) introduces a bi-directional communication primitive called *opportunistic anycast*. This primitive establishes a data exchange with the *first* neighbor to wake up. In this way, SOFA avoids the need for neighborhood discovery and minimizes the inefficient rendezvous time typical of asynchronous MAC protocols.

By selecting opportunistically the next neighbor to communicate with, SOFA provides a perfect building block for gossiping algorithms [2, 14, 13]. Gossiping techniques have been proven to be particularly suitable to disseminate information in large-scale distributed systems. Overall, SOFA offers an alternative to the traditional protocol stack, that cannot operate in extreme sensor networks.

We implemented SOFA in Contiki and evaluated it through a series of experiments on a 100-node testbed, and with Cooja simulations (considering mobility and densities of up to 450 nodes). Overall, our study makes the following three key contributions:

- We introduce EWSN and the problem of communicating in such networks.
- We present the design, implementation and evaluation of SOFA, a communication protocol that utilizes *opportunistic anycast* to overcome the limitations of inefficient rendezvous mechanisms. SOFA combines the energy efficiency typical of low-power MAC protocols with the robustness and versatility of gossip-like communication.
- We show that SOFA can successfully deliver messages, regardless of mobility, in networks with densities of hundreds of nodes while maintaining the duty cycle at approximately 2%.

2 Related work

The constrained energy resources of WSN led to a first generation of protocols that traded bandwidth utilization for lower energy consumption. Such protocols are based on asynchronous radio duty-cycling methods, which implies that senders need to wait for their receiver to wake up (to rendezvous), before sending their data. While in low power listening (LPL) [3] nodes send a beaconing sequence until the receiver wakes up, in low power probing (LPP) [5, 16], the sender waits for a wakeup beacon from the receiver.

The WSN community is well aware of the limitations of the first generation of low-power protocols and several notable contributions have improved their performance. To reduce the overhead of the rendezvous phase, protocols such as WiseMAC [7] keep track of the wakeup periods of their neighbors and use this information to wakeup just a few instants before the intended receiver. This type of protocols works very well on stable networks, where the overhead of estimating the wake periods is seldom done. Highly mobile scenarios, however, prevent the use of these methods.

Another efficient way to disseminate information has been recently proposed by Ferrari et al. [8]. By using a finely synchronized TDMA mechanism, together with extremely efficient network-wide floods, the authors are able to disseminate information irrespective of mobility. However, their low-power wireless bus requires a central coordinator and a network with limited diameter (the synchronization degrades as the number of hops increases). The wide scale (graph diameter) of EWSN prevents the use of this approach, and the central coordinator exposes a single point of failure.

2.1 The need for opportunistic communication

Several notable studies have identified the important role that opportunistic communication has on improving the performance of low-power WSN. In essence, the key idea of these studies is the following: instead of waiting for a predefined node to wake up, opportunistically transmit to who is available now. In ORW [10], the authors propose to use anycast communication to improve the performance of CTP [9], the de-facto data collection protocol in WSN. In Backcast [6], the authors show that by using anycast communication, the capture effect can be leveraged to increase the probability of receiving an ack from a viable receiver. While SOFA is motivated and inspired by these studies, there is an important difference. We do not use opportunistic anycast to improve the performance of traditional network protocols under mild conditions, but to enable a new communication protocol that scales to EWSN. To summarize we make the following observations:

- Basic asynchronous low-power MAC protocols collapse under the densities and traffic demands of extreme wireless sensor networks.
- Mobility prevents the use of more complex and efficient low-power MACs.
- Opportunistic behaviors are needed to scale to EWSN.

3 SOFA Mechanism

The design of SOFA follows two main goals: reduce the inefficient rendezvous phase of low-power MAC protocols, and guarantee that the dissemination of

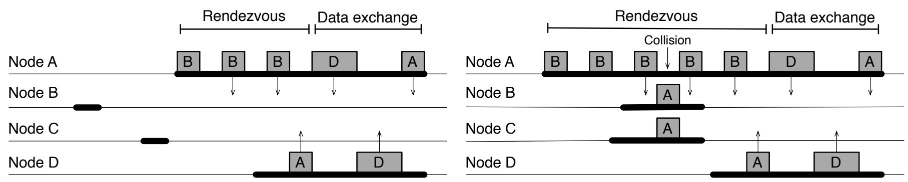

(a) Normal conditions (b) Collision resolution (cf. Section 4)

Fig. 1: SOFA mechanism

data is performed in an efficient and reliable way. To satisfy these goals, SOFA implements an efficient communication primitive, called *opportunistic anycast*, that minimizes the rendezvous overhead and natively supports Gossip, a robust data dissemination technique created for large-scale networks.

Before proceeding it is important to remark that SOFA focuses on maximizing the messages exchanged locally among neighbors (1-hop), leaving the multi-hop dissemination and aggregation of information to the Gossip layer.

3.1 The basic idea

The general idea of SOFA can be applied to any asynchronous duty cycled MAC protocol. Due to space constraints, we focus our analysis on the LPL version of SOFA. The reason is that this implementation performs better in extreme densities, especially in terms of reliability. Nevertheless, in Appendix A, we will provide some insight on the LPP implementation.

Rendezvous phase: In traditional LPL protocols [3], when a sender wakes up, it transmits a series of short packets –called beacons– and waits for the receiver to wake up. When the intended receiver wakes up, it hears the latest beacon and sends an acknowledgement back. SOFA follows a similar mechanism: the sender, node A in Figure 1(a), also broadcasts a series of beacons but only waits until *any* neighbor wakes up. The main difference between the two mechanisms lays in the selection of the destination. While in LPL the destination is chosen by the upper layers in the stack, in SOFA the MAC protocol opportunistically chooses the destination that is most efficient to reach: the first neighbor to wake up. If nodes B or C were to be chosen, node A would need to send beacons (jam the channel) until these nodes wake up again. By sending its data to the first neighbor that wakes up (node D), SOFA reduces the nodes' rendezvous time, allowing low-power MAC protocols to efficiently scale to EWSN. We call this communication primitive *opportunistic anycast*.

Data exchange phase: Selecting the first (random) neighbor that wakes up as *the* destination, has a strong relation with a family of randomized networking algorithms called gossiping [2, 14]. *Gossip algorithms do not aim for traditional end-to-end communication* (where routes are formed and maintained ahead of time), instead they exchange information randomly with a neighbor (or subset of neighbors). The relation between SOFA and Gossiping is fundamental for the practical impact of our work. Unicast and broadcast primitives allow the development of a wide-range of algorithms and applications in WSN such as routing, data collection, querying and clustering (to name a few). Unfortunately, under the stringent characteristics of EWSN these basic primitives collapse. Our aim is to provide an alternative communication protocol for extreme conditions. We hope that this effort will allow the community to use SOFA as a basic building block for other gossip applications such as routing in delay tolerant networks [15] and landscaping of data [11].

We will now describe the design of the three key characteristics of SOFA: *short rendezvous phase, reliable push-pull data exchange,* and *random peer sampling.* The design of a short rendezvous phase was influenced by the limitations of asynchronous duty cycled protocols. The push-pull data exchange and the random peer sampling were designed to satisfy the needs of general gossiping applications.

3.2 Short rendezvous phase

Stopping at the first encounter, instead of searching for a specific destination, has two important consequences on the performance of SOFA. First, and most importantly, it eliminates the main limitation that LPL has under extreme networking conditions: channel inefficiency. By drastically reducing the length of the rendezvous phase, the channel no longer gets easily saturated by medium/high traffic demands or medium/high node densities. A short rendezvous phase also reduces the duty cycle of the radio, which in turn, increases the lifetime of the node. Second, increasing the network's density (up to a point) improves the performance of SOFA. With more neighbors, the probability that one will soon wake up is higher.

To quantify the benefits of a short rendezvous phase, we present a simple model that captures the expected duration of the rendezvous phase as a function of the neighborhood size and the wakeup period (the time elapsed between two consecutive wake-ups of a node). Since nodes wake up periodically in a completely desynchronized way, we can model the inter-arrival times of the nodes' wake-ups as a set of independent random variables with uniform distribution. The first order statistic U_1 can then be used to estimate the length of the rendezvous phase. The expected length $E[U_1]$ of N uniform random variables (neighbors) is given by the Beta random variable with parameters α=1 and β=N

$$U_1 \sim B(1, N), \quad E[U_1] = \frac{1}{1+N}$$

Given a wakeup period W and a neighborhood size N, the expected length of the rendezvous phase of SOFA can be computed as follows:

$$E[s] = \frac{W}{1+N} \tag{1}$$

Considering that the expected rendezvous time of unicast $E[u]$ is $W/2$ [3, 12] and that the time spent for the data exchange phase is negligible compared to

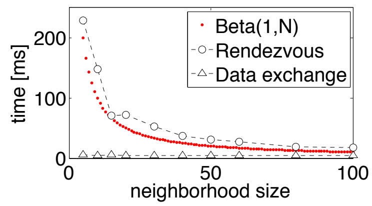

(a) Rendezvous time of SOFA compared to the Beta model

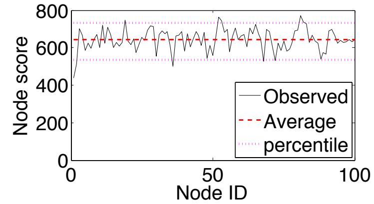

(b) Node selection. The score shows how many times each node is selected (cf. Section 3.4)

Fig. 2: SOFA rendezvous phase (testbed results)

the rendezvous (see Figure 2(a)), we can model the gain G of SOFA compared to unicast as the following:

$$G = \frac{E[s]}{E[u]} = \frac{W}{1+N}\frac{2}{W} = \frac{2}{1+N}$$

For a node with 99 neighbors, this means that the expected rendezvous times of SOFA is 50 times smaller than the one using unicast. Figure 2(a) compares the expected length of the rendezvous phase using the proposed model with values observed in testbed experiments. In this example, W=1 s and the neighborhood size ranges from 5 to 100 nodes. The slight underestimation is mainly due to collisions, which delay the detection of the first node by the sender.

It is important to highlight three key points about the impact of density on SOFA. First, since the performance of SOFA is not significantly affected by changes in medium/high densities, SOFA does not need to adapt to this type of density fluctuations in mobile networks. Second, to reduce the duration of the rendezvous phase in low density networks, the wakeup period can be reduced (at the cost of increasing the duty cycle). This trade-off is studied in more detail in Section 5.3. Finally, in case the network switches from an extreme condition to a normal one (low density), the protocol stack can switch to the use of standard broadcast and unicast messages. To detect the density of the network, SOFA can exploit the tight correlation between the number of neighbors and the expected length of the rendezvous phase (Equation 1).

3.3 Reliable push-pull data exchange

To exchange data efficiently and reliably, SOFA has two phases: a 2-way rendezvous phase and a 3-way data exchange phase. These phases are shown in Figure 1(a) and their design is driven by two factors: (i) the high relative cost of the rendezvous phase compared to the data-exchange phase, and (ii) the effect of unreliable and asymmetric links on the *constant mass* requirement of gossip's *data-aggregation* algorithms. The effect of these factors is explained below.

Using a push-pull exchange amortizes the high relative cost of the rendezvous phase. Gossiping algorithms have two types of data communication:

push and push-pull. In the push method, only the sender transfers information to the receiver(s). In the push-pull method, two nodes exchange their information. Compared to the latter, push-pull allows gossip algorithms to compute more complex aggregates and converge faster [4]. Nevertheless, from our perspective what matters most is the relative cost of the rendezvous phase. Given that the cost of this phase is high compared to the data exchange phase, it is beneficial to exchange as much information as possible once two nodes rendezvous. For this reason, SOFA implements a push-pull approach. A push approach would double the overhead of the rendezvous phase, making SOFA less resilient to extreme conditions.

The 2-way rendezvous phase filters out asymmetric and unreliable links, while the 3-way handshake reduces the probability of loosing "gossip mass". Loosing messages has a particularly detrimental effect on the accuracy of gossiping. For example, when two nodes agree to swap half their value (mass), the loss of a message results in a too low value on the node that missed it, which influences the outcome of all the other nodes as the error propagates in consecutive rounds. The conservation of mass is, thus, an important issue in gossiping algorithms. From a design perspective, this means that we need to consider two important points. First, nodes should avoid the use of unreliable and asymmetric links (which have been shown to be commonplace in WSN [18]). Second, if a packet is lost, we have to reduce the chances of loosing mass.

The 2-way rendezvous phase reduces the chance of using unreliable and asymmetric links. Several studies have shown that unreliable links are usually asymmetric (and vice versa) [18]. On the other hand, bidirectional links are usually characterized by being more reliable. By performing a 2-way exchange before transmitting the actual data, SOFA increases the chances of using a reliable link. It is important to remark that some LPL methods do not follow this approach [12]. These methods piggyback the data on the beacons and acknowledgement packets, that is, they transmit information without checking first if the link is reliable and symmetric or not.

The 3-way data exchange phase reduces the chance of loosing mass in the event that a packet is lost. In spite of our efforts to filter out unreliable and asymmetric links during the rendezvous phase, the high temporal variability of low-power links can cause a reliable link to become momentarily unreliable. In the event that a packet is lost, the worst situation for two nodes is to disagree on the outcome of an event. That is, two nodes should either agree that the message exchange was successful (both nodes received the mass) or agree that no message was exchanged (aborting the exchange). If only one node deems the event as successful, then the mass of the other is lost. The latter situation happens when the last packet of an n-way handshake is lost. This (dis)agreement problem is discussed in depth in [1], and the authors prove that in WSN the best strategy to reduce disagreements is to use a 3-way handshake.

3.4 Random Peer sampling

Most gossip algorithms rely on the selection of a *random* neighbor (or subset of neighbors) at each round. Having a good random selection leads to a faster convergence. To ensure a proper random selection, SOFA introduces random values to the wakeup periods of each node. For a wakeup period of W seconds, nodes wake up uniformly at random between [$0.5W$, $1.5W$].

To validate the effectiveness of our approach, we performed an experiment on a 100-node testbed. For 10 minutes nodes exchange messages and count the number of times they are selected by their neighbors (their *score*). Figure 2(b) shows that the distribution of the scores is close to uniform, with the [5, 95] percentiles close to the average value. It is important to remark that this evaluation was performed on a static testbed. Mobility would further randomize the selection of neighbors, facilitating the dissemination of data, and drastically reducing the convergence time of Gossip [13].

4 Implementation

We implemented SOFA on the Contiki OS based on X-MAC [3]. Nodes were configured to wakeup every second for 10 ms. If a beacon is received within this 10 ms period, the node sends an acknowledgement and starts the data exchange phase. Otherwise, the node goes back to sleep. Notice that these parameters set a minimum duty cycle of 1 %, hence, any extra activity beyond this point is part of the overhead caused by the rendezvous and data exchange phases. Below we describe the implementation of the most important features of SOFA.

Transmit back-off. In traditional MAC protocols, before sending a packet, a transmitter first checks the signal level of the channel (CCA check) to see if there is any activity. If no activity is detected the packet is sent. In SOFA, we do not perform a CCA check. Instead, a potential sender listens to the channel for 10 ms acting, practically, as a receiver. If after this period no packet is detected, the node starts the rendezvous phase. If the node detects a packet that is part of an on-going data exchange, it goes back to sleep (collision avoidance). However, if the detected packet is a beacon, the node changes its role from sender to receiver. By performing a transmit back-off instead of a CCA check, *SOFA transforms a possible collision between two senders into a successful message exchange with a very low rendezvous cost.*

Collision avoidance. One of the key challenges of operating under extreme density conditions is the higher likelihood of collisions due to higher traffic demands. SOFA follows a simple guideline to reduce the frequency of collisions: if a sender detects a packet loss –for instance, by not receiving an ack–, instead of attempting a retransmission, the node goes back to sleep. This *conservative approach* reduces the traffic in highly dense networks. The main caveat of this approach is when the lost packet is the last data ack. In this case, the two parties will disagree on the data delivery, causing an information (*mass*) loss. Fortunately, our testbed results show that this is not a frequent event.

There is a collision event that is not avoided by the above mentioned approach and has a higher probability of occurrence in SOFA. When two or more active receivers detect a beacon, their ACKs are likely to collide (cf. nodes B and C in Figure 1(b)). The sender will receive neither of the ACKs and will consequently continue transmitting beacons. Upon receiving a subsequent beacon (not the expected data packet), the two colliding receivers infer that a collision has occurred and both will go back to sleep. The first node to wake up after the collision (node D) will acknowledge the beacon and exchanges its data. Finally, randomizing the wake up periods of nodes helps in reducing the chances that this type of collisions occurs repeatedly among the same couples of nodes.

Packet addressing. SOFA uses two main types of data packets. For the rendezvous phase, the beacons have no destination address, any node can receive and process the information. For the data exchange phase, the packets contain the destination address of the involved parties. The beacon packets were as small as possible (IEEE 802.15.4 header + 1 byte to define the packet type and 1 byte when addressing is needed).

5 Results

To evaluate the effectiveness of SOFA we ran an extensive set of experiments and simulations. Our testbed has 108 nodes installed above the ceiling tiles of an office building. The wireless nodes are equipped with a MSP430 micro-controller and a CC1101 radio chip. To reach the highest possible neighborhood size, we set the transmission power to +10 dBm. With these settings, the network is almost a clique. For our simulations we used Cooja, the standard simulator for Contiki. We tested network densities of up to 450 nodes and different mobility patterns. Simulations beyond this density value are not possible with normal cluster computing facilities. For both experiments and simulations, the baseline scenario was configured to have a wake up period W=1 s and a transmission period T=2 s. That is, nodes wake up every second to act as receivers and every two seconds to act as senders. Considering that nodes listen for packet activity at each wakeup for 10 ms, the baseline duty-cycle is $\approx 1\,\%$. Any extra consumption beyond 1 % is caused by SOFA. The evaluation results presented in this section consider also other values for W and T, but unless stated otherwise the experiments are carried out using the baseline parameters. The results are averaged over 20 runs of 10 minutes each.

5.1 Performance metrics

The evaluation of SOFA focuses on three key areas: energy consumption, bandwidth utilization and *mass* conservation. To capture the performance of SOFA in these areas, we utilize the following metrics:

Duty cycle. The percentage of time that the radio is active. Duty-cycle is a widely utilized proxy for energy consumption in WSN because radio activity accounts for most of the energy consumption in WSN nodes.

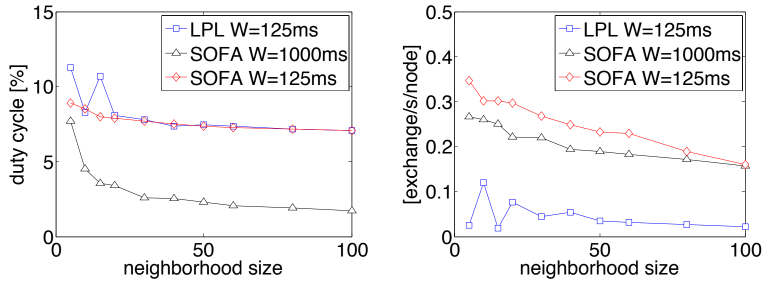

Fig. 3: SOFA compared to LPL (testbed results).

Exchange rate. The number of successful data-exchanges (3-way handshakes) in a second. This is a per-node metric. If, instead, we count the total number of data exchanges in a second over the entire network we refer to the *global exchange rate*.

Mass delivery ratio. The percentage of times that the data-exchange phase ends up without any information loss. Recall that if the ack of the data phase is lost, the receiver deems the exchange as successful, but the sender deems the exchange as unsuccessful and ignores the previously received packet (mass loss). The *mass delivery ratio* is a metric focused on evaluating the viability of SOFA as a basic communication primitive to general gossip algorithms.

5.2 The need for a novel approach

Previously in this paper, we argued that traditional low-power methods collapse under the stress imposed by extreme networking conditions. This subsection quantifies this claim. We compare SOFA with the standard Contiki implementation of LPL on our testbed. To provide a fair comparison, LPL chooses a random neighbor from a pre-computed list of destinations at every transmission request. That is, we do not enforce on LPL the necessary neighbor discovery process that would be needed to obtain the destination address (SOFA does not need an address to bootstrap the communication).

Figure 3 compares the duty cycle and the exchange rate of SOFA and LPL in our testbed. For LPL, the evaluation shows only the result for $W{=}125\,ms$ because LPL collapses with the baseline $W{=}1\,s$. This collapsing occurs because, with $W{=}1\,s$, the rendezvous phase of LPL requires on average $0.5\,s$. Hence, 5 nodes require on average a 2.5-seconds window to transmit their data, but the transmission period is $2\,s$, which leads to channel saturation. Comparing the best parameter for LPL ($W{=}125\,ms$) with the best parameter for SOFA ($W{=}1\,s$) shows that SOFA widely outperforms LPL. For most neighborhood sizes (30 and above), SOFA uses four times less energy and delivers five times more packets for the same T.

It is important to remark that SOFA is not a substitute for traditional low power methods, as they aim at providing different services. SOFA cannot provide several of the functionalities required by applications relying on unicast and broadcast primitives. Most WSN applications are designed for data gathering

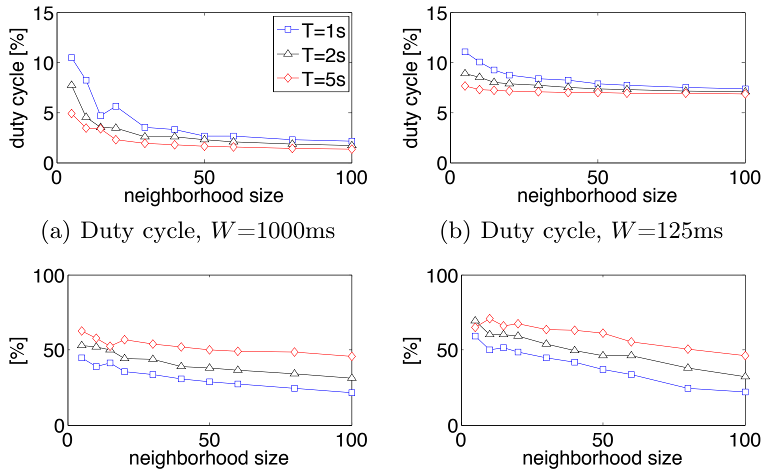

(a) Duty cycle, W=1000ms (b) Duty cycle, W=125ms

(c) Exchange rate, W=1000ms (d) Exchange rate, W=125ms

Fig. 4: Testbed performance for different wake-up times W and transmission periods T. Note that the *exchange rate* is normalized to T.

applications sending a few packets per minute. In these scenarios, the state-of-the-art in WSN research performs remarkably well. The aim of our comparison is to highlight that traditional methods were not designed to operate under extreme conditions neither to efficiently support Gossip applications. We will now analyze the performance of SOFA based on different parameters and scenarios.

5.3 Exploring SOFA parameters

SOFA is a simple protocol with only two parameters available for fine-tuning: the wakeup period W and the transmission period T. We now evaluate the performance of SOFA as a function of these parameters. The results of this subsection are all based on testbed experiments. Figure 4 shows the performance of SOFA for two different wakeup periods (125 and 1000 ms), and for three different transmission periods (1, 2 and 5 seconds).

The impact of the transmission period T. Let us start by analyzing the impact of T on the duty cycle. Figure 4(a) shows two important trends. First, beyond a certain neighborhood size (≈ 30), T does not have a significant impact on the duty cycle. Decreasing the transmission period certainly increases the duty cycle of the node, but not by much. Second, in low/medium dense networks (below 30 neighbors), increasing T has a more significant effect on the duty cycle, but it is still a sub-linear relation. An increment of T by a factor of five, increases the duty cycle by only a factor of two. The reason for the difference in duty cycle between low/medium and high density networks, is that at lower densities, SOFA spends more time on the rendezvous phase. This implies a higher overhead at each transmission attempt. Conversely, increasing the density increases the likelihood of finding a receiver sooner.

Note that, thanks to the *transmit back-off mechanism* (which changes the role of senders to receivers to reduce collisions), increasing the transmission rate decreases the length of the rendezvous phase. With nodes sending data more often, the probability that two senders are active at the same time is higher. While in a normal MAC protocol this would lead to collisions, in SOFA it translates into an efficient message exchange (the rendezvous time is minimal) among the two senders. As for the impact of T on the relative exchange rate, SOFA behaves as most protocols do when they work under high traffic demands: the higher the traffic rate, the more saturated the channel, and the lower the probability to exchange information. This trend is observed in Figure 4(c). It is important to notice, however, that the exchange rate decreases in a gentle manner.

The impact of the wake period period W. Intuitively, reducing the wakeup period should reduce the rendezvous time (because nodes wakeup more frequently), which in turn should free up bandwidth and allow a higher exchange rate. However, the trade-off for a more efficient use of bandwidth would be a higher duty cycle. Figures 4(b) and (d) show the performance of SOFA with a wakeup period W=125 ms. With this value, the baseline duty cycle is 8 %. The figures show that reducing W does increase the relative exchange rate, but mainly on low/medium dense networks (by $\approx$50%). Therefore, it is possible to improve the performance of SOFA in low density networks at the cost of increasing the energy consumption. For high density networks, however, we have a similar throughput but with a duty cycle that is four times higher.

5.4 SOFA under extreme densities and in mobile scenarios

The previous testbed results show that SOFA performs well in densities as high as 100 neighbors. However, from a practical perspective it is important to determine (i) the saturation point of SOFA, i.e., how many nodes SOFA can handle before saturating the bandwidth, and (ii) the impact of mobility. Unfortunately, there are no large-scale mobile testbeds available in the community, and hence, we rely on the Cooja simulator to investigate these aspects.

SOFA shows a strong resilience to extreme densities. Figures 5(a) and (c) show the prior testbed results together with the simulation results. These results consider clique networks for both the testbed and simulation results. First, it is important to notice that Cooja captures, in a pretty accurate way, the trends observed on the testbed. Figure 5(a) shows that the duty cycle continues to decrease (almost monotonically) and stabilize after a density of more than 100 neighbors. Figure 5(c) shows that the exchange rate degrades monotonically but in a graceful manner (notice that the x-axis is in a log scale).

There is, however, a more important question to answer about SOFA: *at what density does it saturate?* The clique curves (bottom two curves) in Figure 5(d) provide some insight into this question. In these experiments, we evaluated the global exchange rate at different densities. For the tested parameters, SOFA saturates when the density approaches 200 neighbors per node. Note that these are clique scenarios. In multi-hop networks, SOFA can exploit the well known

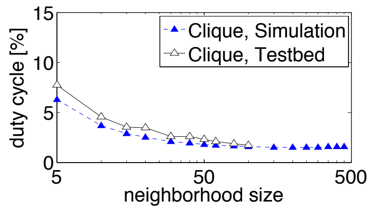

(a) Duty cycle

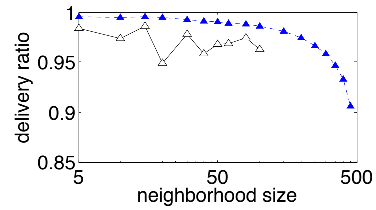

(b) Mass delivery ratio

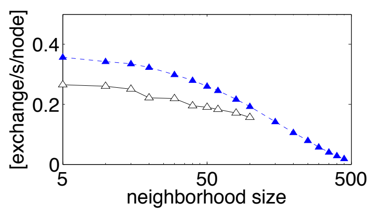

(c) Exchange rate

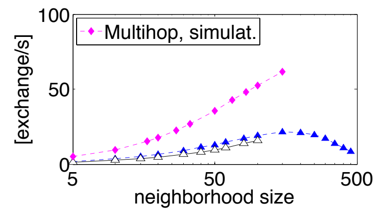

(d) Global exchange rate

Fig. 5: SOFA's performance in extreme network conditions (testbed and simulation results).

spatial multiplexing effect (parallel data exchanges) and achieve higher global exchange rates. The top curve in 5(d) depicts this behavior. The highest point represents a network with 450 nodes and an average density of 150 neighbors.

The performance of SOFA remains the same in static and mobile scenarios. By being a stateless protocol, with nodes acting independently in an asynchronous and distributed fashion, SOFA does not require spending energy on maintaining information about the node's neighborhood and it is independent from the network topology and mobility.

To test SOFA with dynamic topologies, we simulated an area of 150x150 meters where nodes moved according to traces generated by the BonnMotion's random waypoint model. We tested three speeds: 0 m/s (static), 1.5 m/s (walking) and 7 m/s (biking). The radio range was set to 50 meters, with every node being connected, on average, to one third of the network. The maximum density was 150 nodes in a 450-node network. The resulting multi-hop networks had an effective diameter of just below three hops, which ensures that hidden terminal effects are taken into consideration. Figure 6 shows the duty cycle and the exchange rate of SOFA under the patterns *static*, *walking* and *biking*. We can see that the speed of the nodes does not influence neither the energy consumption nor the delivery ratio and the exchange rate.

5.5 Gossip support

As mentioned before, one the goals of our study is to develop a communication primitive that is useful for general gossip applications. In gossip, it is important to conserve mass. Our 3-way handshake phase guarantees that, unless the last ack is lost, the two nodes will reach either a positive agreement (both nodes

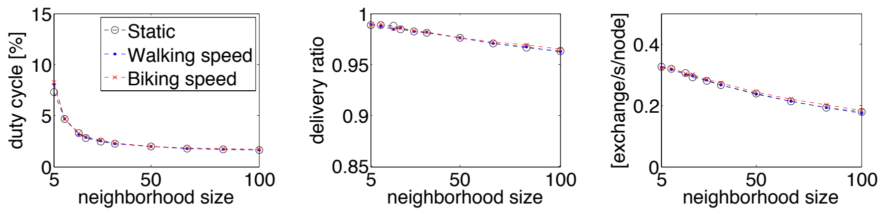

(a) Duty cycle (b) Mass delivery ratio (c) Exchange rate

Fig. 6: SOFA performance under different mobility scenarios (simulation results).

exchange their mass) or a negative agreement (both nodes keep their mass). Clearly, a positive agreement is the most desirable outcome, but both outcomes guarantee that no mass is lost during the exchange. The most important issue is to reduce the possibility of disagreements (when only one node, the sender, deems the transaction as successful).

To evaluate SOFA's ability for mass conservation, we compute the mass delivery ratio. This metric represents the fraction of data exchanges that end up successfully. Figure 5(b) depicts the *mass delivery ratio* of SOFA under different densities. The figure shows that even under extreme densities (450 neighbors) SOFA is able to achieve a high percentage of successful exchanges (above 90 %). This is an important result. In the previous subsection, we found that SOFA saturates at approximately 200 nodes, beyond this point the exchange rate decreases monotonically. But, Figure 5(b) shows that the few exchanges that are able to occur beyond this point are able to be completed successfully. In other words, even under extremely demanding conditions SOFA has a remarkable ability to conserve mass. This feat is due to the careful design of SOFA aimed at (i) selecting reliable links (rendezvous phase), (ii) implementing a transmit backoff instead of a CCA (to avoid sender-based collisions), (iii) avoiding the use of retransmissions (which would jam the channel) and (iv) providing a method to reduce mass losses due to packet drops (3-way handshake).

6 Conclusions

In this study we define the concept of Extreme Wireless Sensor Networks and propose SOFA, a communication protocol that can operate efficiently and reliably under EWSN's stringent conditions. To the best of our knowledge, this is the first effort aiming at mobile sensor networks with high densities and high traffic rates. Traditional low-power protocols, which were not designed with these requirements in mind, simply collapse under such circumstances.

SOFA shows a strong resilience to extreme networking conditions. First, the stateless, asynchronous and distributed characteristics of SOFA make it immune to mobility. SOFA has the same performance in static and mobile environments. Second, SOFA reaches a bandwidth saturation at densities close to 200 nodes and it is able to provide reliable communication for densities of up to 450 nodes.

Finally, it is important to stress that SOFA is not intended to replace traditional low-power methods, as these provide a different set of services. SOFA complements these traditional methods. If the network's conditions change (to a milder state), the network stack can switch to unicast and broadcast primitives.

References

1. C. Boano and M. Zuniga. JAG: Reliable and Predictable Wireless Agreement under External Radio Interference. In *IEEE 33rd Real-Time Systems Symposium (RTSS)*, pages 315–326, 2012.
2. S. Boyd and A. Ghosh. Gossip algorithms: Design, analysis and applications. In *IFOCOM 2005. 24th Annual Joint Conf. of the IEEE Computer and Communications Societies*, pages 1653–1664, 2005.
3. M. Buettner, G. Yee, and E. Anderson. X-MAC: a short preamble MAC protocol for duty-cycled wireless sensor networks. In *Sensys '06 4th int. conf. on Embedded networked sensor systems*, pages 307–320, 2006.
4. A. Demers, D. Greene, and C. Hauser. Epidemic algorithms for replicated database maintenance. In *Sixth ACM Symposium on Principles of distributed computing*, pages 1–12, 1987.
5. P. Dutta, S. Dawson-Haggerty, Y. Chen, C.-J. M. Liang, and A. Terzis. Design and evaluation of a versatile and efficient receiver-initiated link layer for low-power wireless. In *Sensys '10 8th ACM Conference on Embedded Networked Sensor Systems*, pages 1–14. ACM Press, 2010.
6. P. Dutta, R. Musaloiu-E, I. Stoica, and A. Terzis. Wireless ACK collisions not considered harmful. In *In HotNets-VII The Seventh Workshop on Hot Topics in Networks*, 2008.
7. A. El-Hoiydi. WiseMAC: an ultra low power MAC protocol for the downlink of infrastructure wireless sensor networks. In *ISCC '04 Ninth Int. Symp. on Computers and Communications*, pages 244–251, 2004.
8. F. Ferrari, M. Zimmerling, L. Thiele, and L. Mottola. The low-power wireless bus. In *10th Conf. on Embedded networked sensor systems (Sensys)*, pages 1–14, 2012.
9. O. Gnawali, R. Fonseca, and K. Jamieson. Collection tree protocol. In *Sensys '09 7th ACM Conf. on Embedded Networked Sensor Systems*, pages 1–14, 2009.
10. O. Landsiedel, E. Ghadimi, S. Duquennoy, and M. Johansson. Low Power , Low Delay : Opportunistic Routing meets Duty Cycling. In *IPSN '12 11th int. conf. on Information Processing in Sensor Networks*, pages 185–196, 2012.
11. A. Loukas, M. Zuniga, M. Woehrle, M. Cattani, and K. Langendoen. Think globally, act locally: on the reshaping of information landscapes. In *IPSN '12 12th Int. Conf. on Information Processing in Sensor Networks*, pages 265–276, 2013.
12. D. Moss and P. Levis. Box-MACs: Exploiting physical and link layer boundaries in low-power networking. *Computer Systems Laboratory Stanford University*, 2008.
13. A. D. Sarwate and A. G. Dimakis. The Impact of Mobility on Gossip Algorithms. *IEEE Transactions on Information Theory*, 58(3):1731–1742, Mar. 2012.
14. D. Shah. Gossip algorithms. *Found. and Trends in Networking*, pages 1–125, 2009.
15. T. Spyropoulos, K. Psounis, and C. S. Raghavendra. Spray and Wait : An Efficient Routing Scheme for intermittently connected mobile networks. In *SIGCOMM 2005 workshop on Delay-tolerant networking*, pages 252–259, 2005.
16. Y. Sun and O. Gurewitz. RI-MAC: a receiver-initiated asynchronous duty cycle MAC protocol for dynamic traffic loads in wireless sensor networks. In *SenSys '08 6th ACM conf. on Embedded network sensor systems*, pages 1–14, 2008.

17. L. Tang, Y. Sun, O. Gurewitz, and D. B. Johnson. EM-MAC: a dynamic multi-channel energy-efficient MAC protocol for wireless sensor networks. In *Twelfth Int. Symp. on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 23:1–23:11, 2011.
18. M. Zuniga and B. Krishnamachari. An analysis of unreliability and asymmetry in low-power wireless links. *ACM Trans. on Sensor Networks*, 3(2):7, June 2007.

Appendix A: SOFA and LPP

SOFA is a generic mechanism that can be applied to any low power duty cycled protocol. In this appendix we briefly explore SOFA in combination with low power probing (SOFA-LPP). The main difference with SOFA-LPL resides in the rendezvous mechanism. While in SOFA-LPL the initiator actively probes the channel with beacons, in LPP the initiator passively listens to the channel until a neighbor sends out his wake-up beacon signaling it is ready for communication.

LPP can lead to collisions when two nodes target the same destination, as they will both respond to the wake-up beacon of that node. For unicast this is unlikely (nodes target specific neighbors), but for opportunistic anycast senders in close proximity target the same (first to wake-up) node. Similar to SOFA-LPL's transmit back-off, we can remedy this by having an initiator first act as a receiver (by sending out a beacon of its own) before turning into a sender (passively listening for beacons); this way two initiators engage in an effective communication with each other instead of ending up with a collision.

Figure 7 shows that SOFA-LPP performs worse than SOFA-LPL. In particular, SOFA-LPP suffers in terms of mass delivery ratio and exchange rate. This is largely due to the periodical wakeup beacons, which saturate the channel for medium to highly-dense networks.

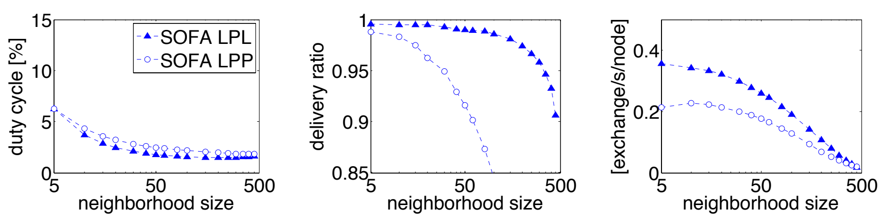

(a) Duty cycle (b) Mass delivery ratio (c) Exchange rate

Fig. 7: SOFA-LPP compared to SOFA-LPL (simulation results).