

How Environmental Factors Impact Outdoor Wireless Sensor Networks: A Case Study

Ramona Marfievici¹, Amy L. Murphy², Gian Pietro Picco¹, Federico Ossi³, Francesca Cagnacci³

¹ University of Trento, Italy, {marfievici,picco}@disi.unitn.it

² Bruno Kessler Foundation, Trento, Italy, murphy@fbk.eu

³ Edmund Mach Foundation, S. Michele all'Adige, Italy, {federico.ossi,francesca.cagnacci}@fmach.it

Abstract—How do the characteristics of the surrounding environment affect the ability of the nodes of a wireless sensor network (WSN) to communicate?

Partial answers to this question can be found in the literature, but always with a focus on the short-term, small-scale behavior of individual links, as this directly informs the design of WSN protocols. In this paper, we are instead concerned with the large-scale behavior of the overall network, observed over a longer time scale, as our primary interest is to support the deployment of WSNs by characterizing the impact of the target environment.

Motivated by a real-world wildlife monitoring application, we report about experimental campaigns in three outdoor environments characterized by varying degrees of vegetation. Experiments are repeated in summer and winter, to account for seasonal variations, and span multiple days, allowing us to assess variations induced by the succession of day and night. Our experiments focus primarily on characterizing the impact of the environment on the physical layer, but we also investigate how this is mirrored at higher layers. We analyze the experimental data along multiple dimensions, yielding quantitative answers to the aforementioned question, and eliciting trends and findings previously not reported in the literature. We argue that this type of study may inspire new methods to better estimate the performance of a WSN in its target deployment environment.

I. INTRODUCTION

The performance of a wireless sensor network (WSN) deployment is obviously affected directly by the quality of the links enabling communication among nodes. In the case of commonly-used IEEE 802.15.4 radios, the link quality is in turn easily affected by a variety of environmental factors, which cause variations over time and space. This is a well-known fact that has been studied extensively in the context of protocol design for WSNs. For instance, a number of techniques exist for estimating the quality of links and determining their best use in MAC and routing protocols [3], [11]. These techniques are themselves based on the findings of studies that empirically determined key properties of low-power wireless communication [9], [12], [18], [21], [22].

The goal of supporting the efficient design of network protocols has led the research community to focus primarily on the study of *small-scale*, short-term variations of the individual links, as this is what matters for determining, say, forwarding decisions in a routing protocol. Further, the need to analyze in detail these variations has led to experimental setups constituted mostly by highly controlled lab settings. Nevertheless, the few exceptions in the literature [5], [7], [15]

show that real-world environments may yield quite different link properties. We review the related work in Section II.

In this paper, we take a different look at the problem of assessing link quality, motivated by deployment and application issues rather than protocol design. Specifically:

- 1) we focus on the *large-scale* changes of the overall network, as induced by different environmental conditions. Therefore, we characterize the *aggregate* constituted by the entire network, rather than the individual links.
- 2) we assess the above *in vivo* (i.e., in an actual real-world environment) in contrast to the *in vitro* (i.e., in labs or controlled setups) experiences reported in the literature.

Application motivation. The original motivation for the work described here stems from a WSN deployment we are pursuing in collaboration with biologists studying the social behavior of wildlife, namely, roe deer in the mountains near Trento, Italy.

The state of the art for these studies is based on GPS devices enabling the tracking of animal movements [6]; social interactions are inferred from the resulting trajectories. Unfortunately, GPS is energy-hungry, hence the sampling rate is typically very low. This fact, along with inaccuracies due to partial sky views, introduces uncertainty and approximation in the trajectories, and therefore on the spatial proximity among individuals inferred from them. Instead, animal-borne WSN nodes enable biologists to record *directly* the interactions among animals. The low-power wireless radio is used as a proximity sensor, and the network-level functionality of neighbor discovery is used directly by the application.

Nevertheless, to be useful to biologists, the interactions recorded by this WSN must be correlated to the characteristics of the environment. Roe deer dwell in habitats with different vegetation, and are tracked year-long, across very different weather conditions. How do the various environmental factors affect the quality of communication? In our application, the answer to this question affects directly the reliability of sensing interactions among animals. However, the answer is of more general interest, given the many applications exploiting a WSN in an outdoor environment. In this context, knowledge about the effect of the environment on communication provides insights about the reliability and lifetime of the WSN.

Contributions. This work reports about the experimental campaigns through which we investigated these issues. Their design was informed by the biologists in our team, interested

in determining how low-power wireless links are affected by:

- *Presence and density of vegetation*: we selected habitat components of a heterogeneous landscape (e.g., mountain range of roe deer): two types of forest, beech (dense vegetation, deciduous) and spruce (sparser vegetation, evergreen), as well as open field.
- *Seasonal variations*: experiments in summer and winter.
- *Daily variations*: experiments span the 24 hours.

The above entail in-field experiments, notoriously effort-demanding and time-consuming. Section III describes TRIDENT, the tool we used in our in-field campaigns. Unlike similar tools in the state of the art, TRIDENT does not require a separate, wired infrastructure, and is expressly designed to simplify the in-field assessment of connectivity, the experiment configuration, and the collection of results.

In a nutshell, our experiments investigate what happens if the *same* WSN is “immersed” in different combinations of the above environmental factors. The topology and experiment execution is the same across the various test environments. We describe the details of our experimental setup in Section IV.

The questions we answer are, for instance: What changes between operating a given WSN in an open field vs. in a forest, or during day vs. night? Some qualitative answers can be derived from existing results, especially if one focuses separately on the impact of each environmental parameter (e.g., foliage, temperature, humidity, snow). Nevertheless, our aim in this paper is to *quantify* the extent of changes, based on the combined effect of the various parameters in different, real outdoor environments. To this end, we rely on the metric of *packet delivery rate (PDR)*, i.e., the ratio of packets received on a link over those sent. However, we aggregate and analyze the *PDR* in different ways and against several combinations of environmental factors to elicit our findings. These are reported in Section V, where we discuss the impact of the environment on the physical layer, and in Section VI where we show how this impact is mirrored at the application layer—the one directly relevant to end users.

The observations we make may inspire new methods to better estimate the performance of a WSN in its target deployment environment, as we argue in Section VII, before the brief concluding remarks of Section VIII.

II. RELATED WORK

Empirical studies of link properties. Many works aim to understand link dynamics (e.g., [9], [12], [18], [21], [22]). Most of them were carried out indoor, in a controlled environment, in contrast with our experiments which are carried out in-field and outdoors. Moreover, these works study the behavior of individual links, usually focusing on short-term variations. In contrast, we are concerned with characterizing the aggregate behavior of the network. Nevertheless, we borrow commonly-used concepts and techniques from these works, e.g., the notion of transitional area and the definition of link classes.

Long-lived outdoor deployments. Several experiences with outdoor WSN exist. Most focus on specific applications,

and only a few report lessons learned in deploying long-lived WSNs. GreenOrbs [14] focuses on large-scale WSN operating in the forests, analyzing network yield and routing performance. Life under your feet [16], a WSN network for soil science, was deployed for more than a year, in two forests where it provided insights about the performance of sensor sampling, energy consumption, and overall software reliability in the field. All these experiences, however, focus on the performance of the final application, and do not aim to characterize the environment where they are deployed and its variations, which is instead our goal in this paper.

Impact of the environment on low-power wireless links.

Several researchers have shown that outdoor WSNs are affected by weather. Thelen et al. [20] described how radio propagation is favored by high humidity in their potato field deployment. Other works [1], [17] suggest that fog and rain may have a severe impact on the transmission range of WSN nodes, especially w.r.t. packet reception. In particular, Boano et al. [5] quantified the impact of these conditions on different platforms in an outdoor industrial setting, showing that light rainfall has a negligible effect on signal strength while heavy rainfall can disrupt connectivity. Bannister et al. [4] showed that high temperature negatively affects communication, based on data from a radio survey in the desert and applied to simulations of localization and data collection.

Many of the findings above are indeed confirmed by our observations. However, in comparison, our work assesses the overall trends induced by the environment in a more systematic and holistic way, by taking into account seasonal and daily variations. Moreover, it also explores the effect of environments with different characteristics (e.g., vegetation).

III. TOOL SUPPORT: TRIDENT

The experimental campaigns we describe in this paper were performed using TRIDENT [10], a tool developed in our group to support the in-field assessment of connectivity. TRIDENT allows untethered execution of experiments, without the wired backbones necessary to other state-of-the-art tools [2], [8], [19]. The latter would be highly impractical in our case, due to the relatively long duration of tests and the environmental conditions (e.g., snow). The tool is easy to use; it has already been used in-field by scientists without technical knowledge about WSNs [7]. TRIDENT supports the entire process, from the design of the field tests to the gathering and sharing of the results, along with support for data visualization and analysis.

Once the tests have been designed, their description is input to TRIDENT, producing the properly configured TinyOS code to be loaded on TMote Sky motes. The tests are run in-field only by battery-powered devices; optionally, a base station can supervise the tests online. Results are retrieved from nodes by TRIDENT via multi-hop forwarding or a direct USB connection, and later uploaded on a shared repository.

The actual probing of communication links is performed by having senders transmit messages in round-robin to avoid collisions, and listeners record packet reception. For each

packet, the sender logs the ambient noise before transmission, and the receiver logs RSSI and LQI values.

TRIDENT allows splitting an experiment into a set of *rounds*, each characterized by a set of parameters. These include the radio channel and power, the number of messages sent per sender (with the indication of how many to send in a single sequence) as well as the time interval between sequences. Each node can be configured to behave as sender, listener, or both.

IV. EXPERIMENTAL SETUP

Time and location of the experiments. The main findings reported in this paper were gathered in two experimental campaigns, performed during the winter (February 21–March 27) and late spring (May 24–May 31) of 2011. A third campaign was performed in the summer (July 9–August 9) of 2012, to validate the findings in a different time period, and explore the impact on the application layer. We describe the differences of this campaign in Section VI. As the weather conditions in late spring and summer campaigns were very similar, hereafter we refer to both as “summer”, for simplicity. Also, weather conditions were uniform within each campaign.

The location of the experiments was chosen to be representative of the environment where our target WSN application is going to be deployed, and to cover different conditions of vegetation. All the experimental sites are on Mount Bondone, near Trento, Italy. Upon suggestion of the biologists on our team, we identified three locations in this area: OPEN is a wide meadow with essentially no trees, while SPRUCE and BEECH are forests characterized by the corresponding tree type.

Hardware/software platform. We used TMote Sky nodes, equipped with the ChipCon 2420 radio chip compliant with IEEE 802.15.4, and on-board omnidirectional antenna. These are popular choices, also allowing us to leverage experience from our previous uses of TRIDENT (e.g., [7]). Each node was placed in an IP65 water-proof box with a transparent cover, containing the two D-size batteries powering the node.

Node placement. The topology of our WSN deployment, shown in Figure 1, consists of 8 nodes arranged in a cross, similar to [7]. This represents a good trade-off between deployment effort and coverage of different link distances, in our case ranging from 7 m to 64 m.

In OPEN, nodes were attached to 3-meter tall wooden poles, planted vertically in the ground. In the two forests, nodes were lashed to trees. We selected the three sites to have a similar terrain inclination. We also ensured the orientation of antennas to be the same in all cases.

Test execution. All nodes were configured both as listeners and senders in TRIDENT. A master (node 0 in Figure 1) disseminated the start time and the transmission schedule for experiment. The latter was divided in 30-minute rounds, during which each of the 8 nodes sent 215 packets, at a

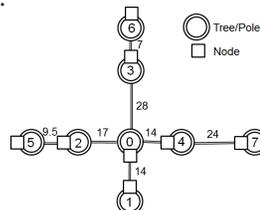


Fig. 1. WSN topology in winter and summer 2011.

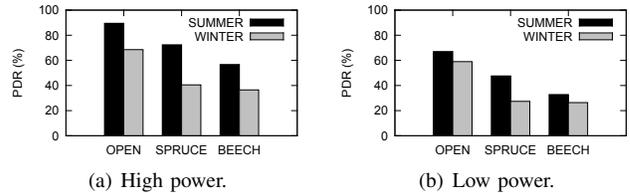


Fig. 2. Network-wide average *PDR*, for all combinations of site (OPEN, SPRUCE, BEECH), season (summer vs. winter) and power (high vs. low).

rate of 1 packet/s. No MAC protocol was used, given our goal of characterizing physical connectivity. In each round, only one node transmits at a time, with a large-enough gap between transmissions to account for the clock drift, ensuring the absence of collisions among senders.

Each node aggregated *PDR*, RSSI, LQI over the 30-minute round. Rounds were executed one after the other without intervention for 2 days, after which they were stopped for one hour (at 9 AM) to download the data from the flash and replace the batteries. Moreover, we interleaved rounds at -1 dBm (power 27, hereafter “high power”) with rounds at -8 dBm (power 14, hereafter “low power”). We used channel 18 in all cases. Overall, the results of 96 rounds of 30 minutes were collected in each environment, for each season.

V. IMPACT OF ENVIRONMENT ON THE PHYSICAL LAYER

The objective of the experiments described here is to understand if and how the environmental conditions, both in terms of morphology (e.g., vegetation present) as well as daily and seasonal variations, affect the low-power wireless links of a WSN, from the standpoint of the physical layer.

A. Network-wide Packet Delivery Rate

We begin our analysis with the roughest indicator of link quality, and yet the most intuitive and directly informative: the *packet delivery rate (PDR)* (i.e., the ratio of packets received over those sent) computed over the entire set of links. Despite its simplicity, this provides an immediate and easy-to-compute macro-indicator telling us how the same network behaves, once immersed in different environments.

Figure 2 shows the results from our experiments. A few trends are clearly identifiable. First, the quality of communication decreases as one progresses from OPEN to SPRUCE to BEECH—i.e., as the quantity of trees and foliage increases. The trend is more marked during summer at both powers: going from OPEN to SPRUCE, and from SPRUCE to BEECH, shows differences in *PDR* of the order of 15-20%. Second, the seasonal variation also induces dramatic changes in link quality, with winter consistently worse than summer, mostly due to the presence of snow. Differences can be as high as 30%, as in the case of SPRUCE at high power. Third, during winter the differences in *PDR* between the two forests, SPRUCE and BEECH, are negligible. As we discuss later, this is not true of other parameters related to link quality. However, through the macroscopic lens provided by the network-wide average *PDR*, it appears like the combination of snow and vegetation yields the same effect regardless of the density of the latter. In this respect, it should be pointed out that, although

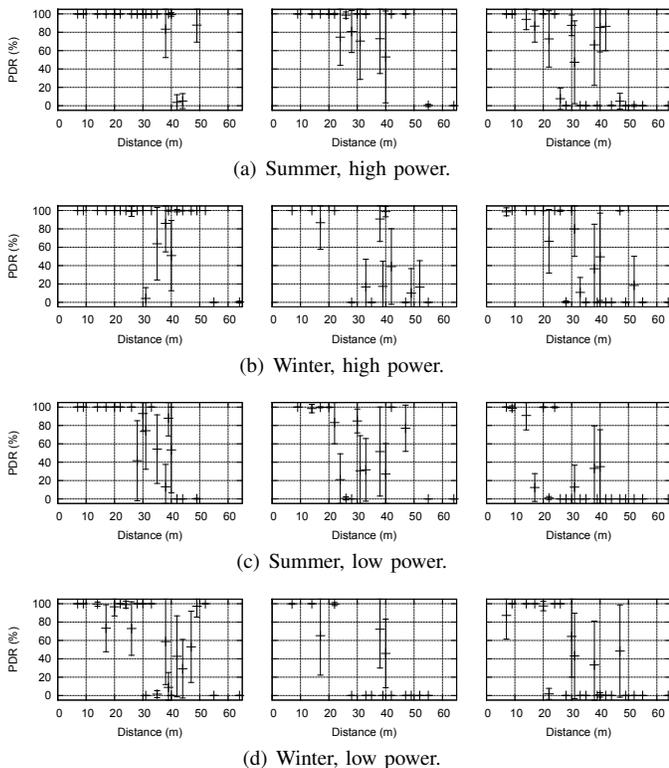


Fig. 3. Transitional areas in OPEN (left), SPRUCE (center), and BEECH (right).

during winter foliage is not present in BEECH (deciduous) while it is in SPRUCE (evergreen), the density of trees and branches is still much higher in the former. Finally, all of the trends above hold both for high and low power although, quite obviously, the PDR values are lower in the latter case.

We dissect further the environment impact on the physical layer, with tools more sophisticated than the average PDR .

B. Transitional region

Several studies [12], [18], [21], [22] have classified low-power wireless links in three distinct reception regions—connected, transitional, and disconnected—based on the distance of the receiver from the sender and on the PDR . The connected region is the closest to the sender: its links are of good quality, and often stable and symmetric. At the other extreme, the disconnected region is the farthest from the sender: it does not contain links practically usable for communication. The transitional region (also referred to as “gray area”) is a mix of the two, and contains links that exhibit a high variance in packet reception rate, as well as asymmetric links. Unfortunately, the transitional region often spans a large fraction of the communication range. The beginning of the transitional region and its span determine how challenging it is to ensure reliable communication in a WSN deployment.

Figure 3 shows the transitional areas for OPEN, SPRUCE, and BEECH, in different seasons and with different power settings, as described in Section IV. Several observations can be made.

OPEN provides the situation most favorable to communication. The transitional area begins (i.e., the connected area ends) much farther than in the forest deployments, for all

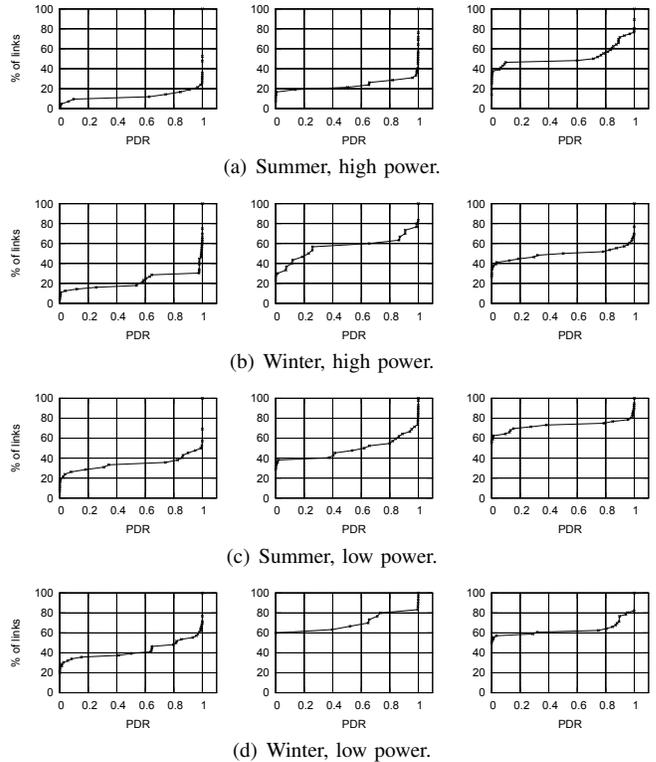


Fig. 4. Cumulative distribution function of PDR in OPEN (left), SPRUCE (center), and BEECH (right).

combinations of seasons and power settings. The presence of vegetation appears to impair communication, by reducing the span of the connected area, and increasing the transitional one.

Seasonal variations are also evident. Indeed, the presence of snow on the ground and on tree branches during winter has a detrimental effect to communication. In OPEN, the transitional area begins closer to the sender at both powers, although the phenomenon is more marked when using low power. Similar considerations hold for the two forests, although the impact on communication of seasonal variations is less dramatic w.r.t. the one of vegetation. In comparison to OPEN, however, the transitional area not only starts earlier, but is also wider.

The two forests are similar in terms of beginning and span of the transitional region. Differences are more marked in summer, when the denser and broader leaves in BEECH yield a smaller connected region, and a wider transitional one.

C. Link classification

Looking at the transitional area allows one to grasp quickly the extent of communication range, but does not yield insights on the *fraction* of links characterized by a given quality.

We provide this “view” in two ways. The first is to compute the cumulative distribution function (CDF) of the links w.r.t. their PDR , as shown in Figure 4. The other is to adopt the classification used in [18], in which a link can be *dead* ($PDR = 0\%$), *poor* ($PDR < 10\%$), *intermediate* ($10\% \leq PDR \leq 90\%$), *good* ($90\% < PDR < 100\%$) or *perfect* ($PDR = 100\%$), as shown in Figure 5.

Figure 4 shows clearly the relative performance of the three environments. As we progress from OPEN, to SPRUCE,

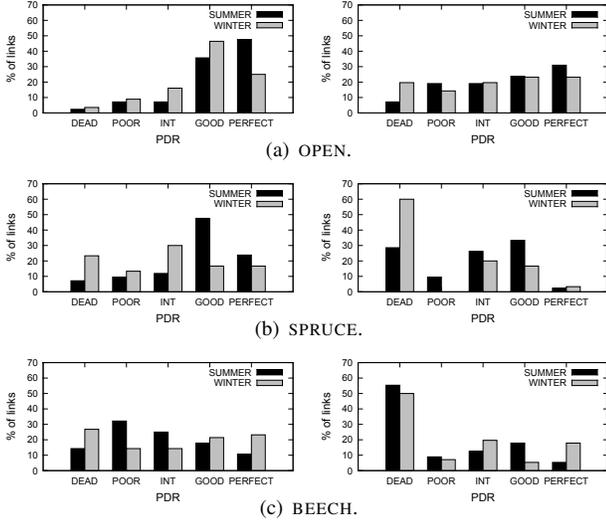


Fig. 5. Seasonal variations, at high power (left) and low power (right).

and finally to BEECH, the CDF “shifts” upwards, due to an increasing fraction of bad (i.e., dead, poor, intermediate) links. This can be seen in Figure 5, where especially the number of dead links increases noticeably as we go from OPEN to BEECH and, dually, the sum of good and perfect links decreases.

Figure 5(a) also shows that winter has a negative impact on OPEN. During summer, at high power, 83% of all links are good or perfect, but only 71% during winter. Further, the number of perfect links drops abruptly from 48% to 25%. This is consistent with our previous considerations about the transitional area “moving” closer to the sender during winter. On the other hand, the impact is much less marked at low power, although the number of dead links doubles.

Similar considerations hold for SPRUCE, in Figure 5(b), although the sum of good and perfect links is smaller than in OPEN, as already mentioned. However, the sum of intermediate, good, and perfect links accounts for 63% in winter and 82% in summer, respectively 40% and 66% with low power.

In BEECH, shown in Figure 5(c), communication at high power in winter is better than in summer, with 44% links being good or perfect during winter compared to 28% during summer, when more intermediate links are also present. The reason for this behavior is the fact that, during winter, foliage is not present in BEECH, and therefore, despite the presence of snow on the dense tree branches, communication is less impaired. A similar trend, albeit less marked, is shown at low power, with many links moving from good to perfect.

The difference between the two forests is less marked in winter: snow reduces the link quality in both SPRUCE and BEECH, but in the latter this fact is partially compensated by the absence of foliage. On the other hand, BEECH is worse in the summer, due to the presence of denser and broader leaves.

D. Day vs. night

We now turn our attention to variations induced by the interleaving of night and day, which affects environmental parameters such as temperature and humidity, which in turn affect link quality. An example is provided by Figure 6,

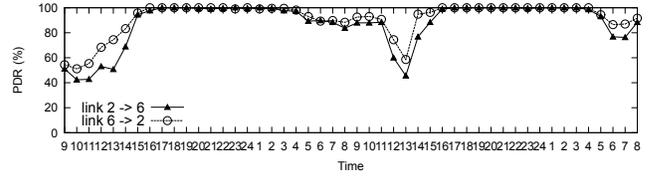


Fig. 6. An example of day vs. night variations: the *PDR* of link 2 → 6 and 6 → 2 in OPEN, during summer, at low power. The node distance is 39 m.

showing the *PDR* between node 2 and 6. The link is perfect at night, but its *PDR* drops significantly (as low as 40%, going from perfect to intermediate) during the day. To analyze in more detail this phenomenon for the entire network, unlike Section V-C we focus only on the charts showing link classification, due to space reasons and to the fact that these better highlight how links “move” across classes when changing from day to night. Figure 7 shows the results of our analysis.

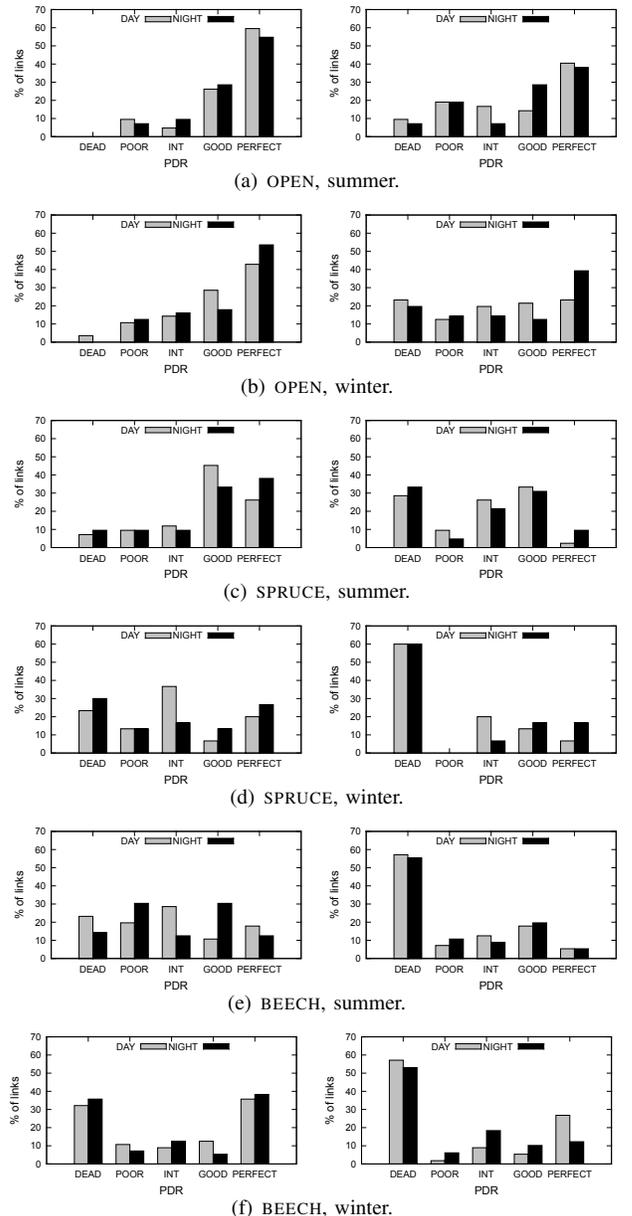


Fig. 7. Day vs. night variations, at high power (left) and low power (right).

In OPEN, variations are very limited during summer, especially at high power. At low power, instead, the number of good and perfect links is higher during the night (66% vs. 54%), as several links move from intermediate to good. During winter, at high power, a significant fraction (10%) of links move from good to perfect when changing from day to night. This trend is even more marked at low power; the number of perfect links doubles and the intermediate links decrease.

In SPRUCE, during summer the number of dead, poor, and intermediate links remains basically unaltered, but perfect links increase by almost 15%. Similar trends are observed at low power, although less dramatic. Interestingly, however, in this case some links move from poor to dead, and others from intermediate to good and perfect. In a sense, the transition from day to night has the effect of “polarizing” the network into dead and perfect links, although overall the variations are somewhat limited. Similar considerations, albeit with more marked variations, hold for winter at high power. At low power, instead, the network is dominated by a significant fraction (60%) of dead links during both day and night, although the shift towards perfect links can still be observed.

As for BEECH, the denser vegetation and broader foliage has the effect of “damping” the daily variations, introducing different trends. During summer, at high power, there is a polarization effect similar to SPRUCE. However, while there is an evident (almost 20%) decrease in intermediate links, there is also a decrease in dead and perfect links. The net effect is that it is the fraction of poor and good links that increases significantly, instead of dead and perfect as in SPRUCE. On the other hand, at low power the change from day to night induces almost no variation, due to the thick foliage that “protects” the links from environmental variations. During winter, thanks to the absence of foliage we observe a polarization towards dead and perfect links as in SPRUCE, although variations are smaller. Finally, at low power the trends we observed in the other cases are reversed: the number of intermediate links doubles, while perfect links are reduced by almost two thirds. We conjecture that this is an effect of the dense branches with snow, creating different “micro-ambients” for communication, although this aspect requires further, finer-grained investigation.

VI. DOES IT MATTER?

FROM THE PHYSICAL LAYER TO THE APPLICATION

In this section, we investigate whether the trends we discussed in Section V bear an impact on the application layer. To this end, we consider two applications. The first is a staple periodic data collection relying on the popular CTP [13] protocol. The second is a controlled deployment of our wildlife monitoring application, where the presence of neighbors in proximity is logged independently at each node. These tests were also the opportunity to verify that the results at the physical layer described in Section V hold in experiments performed at a different time and with a different setup.

A. Periodic data collection with CTP

Experimental setup. These experiments were performed in the summer (July 9–August 9) of 2012. Due to some logistical issues, we performed the tests only in OPEN and SPRUCE. The site for SPRUCE was the same, while the one for OPEN was different but with the same characteristics as the one in Section IV, which meanwhile had become private property.

We also used a different topology w.r.t. Section IV. Indeed, as the goal in this case was to build a tree, we wanted *i)* to have more nodes to build the topology from, and *ii)* to have a uniform node placement to avoid biasing the tree construction.

The topology we used consists of 16 nodes distributed in a grid, as shown in Figure 8. We first deployed it in SPRUCE, where node placement was constrained by distances between trees, and then reproduced it in OPEN with the same distances. As in Section IV, this yielded a rich set of links at distances varying from 10.2 m to 60 m. The nodes are the same as in Section IV, in terms of hardware, packaging, and setup.

We ran two types of experiments on this topology. First, we ran the physical layer ones described in Section IV, using the same hw/sw configuration. The only difference was that, to maintain rounds of 30 minutes with an increased number of nodes, we configured each node to transmit 115 packets per round. This allowed us to verify that the trends observed in Section V held also in the new experiments, as discussed later.

Next, we ran the experiments with data collection. Nodes were configured with the standard settings for CTP and BoXMAC. Nodes acted as sources sending messages towards the sink (node 0) every 30 s. Unlike the physical layer tests, transmissions were not centrally scheduled. We used a packet size of 105 B, including a 94 B payload, representative of many WSN applications. As before, experiments were divided in 30-minute rounds, performed by interleaving high power and low power settings. The collection tree was reset before each round; we ignore the first few minutes of data to account for the tree building phase. A total of 330 30-minute experiments were run in each environment. Unfortunately, when analyzing the data we discovered that a few nodes (2, 4, 6, 14) malfunctioned, and removed them from the data set.

Confirming previous observations at physical layer. When we went back to the field, the first question we wanted to answer was: To what extent the trends we observed in the previous campaign remained the same? The experiments we report about in this section were performed one year later, with a different topology and, in the case of OPEN, in a different albeit similar site. We clearly hoped for close-enough results,

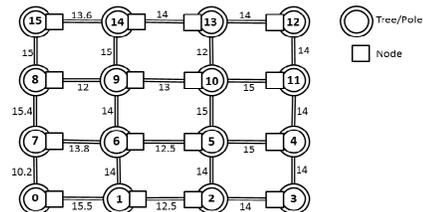


Fig. 8. Network topology for experiments in summer 2012.

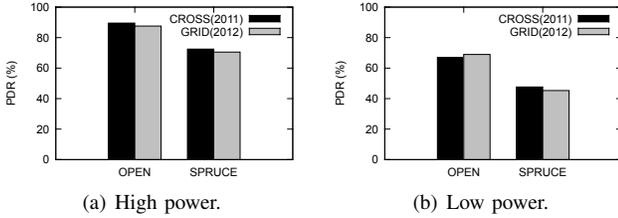


Fig. 9. Network-wide average PDR: comparison of the experiments in Section V and VI, performed one year apart and with different topologies.

but what we distilled from the analysis of the experimental data went beyond our expectations.

Indeed, as Figure 9(a) shows, the difference in network-wide *PDR*, for *all* the combinations of environments we tested, remains within 2–3%. Two considerations are worth making. First, although the WSN topologies we used are different (cross vs. grid) we chose the node distances in the latter to approximate the ones in the former. In other words, the grid covers a set of link distances similar to the cross and, since the grid contains twice the nodes, with more statistical relevance. Second, and most important, the difference in network-wide *PDR* across the two sets of experiments is so small that it suggests that, when one looks at the *aggregate* link quality of the entire network, indeed the impact of the environment on a WSN deployment is relatively stable, on a seasonal scale.

In general, the WSN behavior in the two campaigns is very similar. Space constraints prevent us from reporting all results as in Section V, but the findings are confirmed. For example, Figure 10 shows the *PDR* variation in the new campaign, confirming over the longer 7-day period the observations in Section V-D, i.e., daily variations in *OPEN* affect the WSN only marginally at high power, but significantly at low power.

Impact on application layer. We looked at three metrics to assess the impact of environment-induced differences in connectivity on the operation of CTP: 1) *delivery rate*, i.e., fraction of messages received at the sink over those collectively sent by the sources; 2) *number of CTP beacons* sent; 3) *number of tree parent changes* occurred. The first metric measures the reliability of communication, while the other two measure the overhead necessary to ensure such reliability.

Figure 11 shows that the trends in delivery rate are very different in *OPEN* and *SPRUCE*. The former remains stable at 100%, except during the central day hours, where the delivery rate drops as low as 96% for low power. On the contrary, in *SPRUCE* the range of variation in delivery rate is less dramatic, with drops limited to 98.5%, even for low power. However,

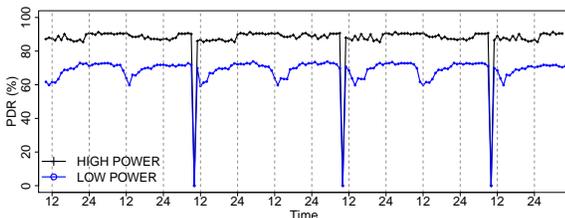
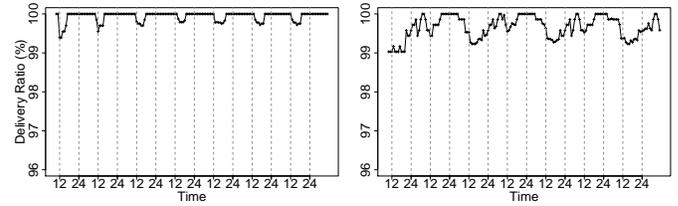
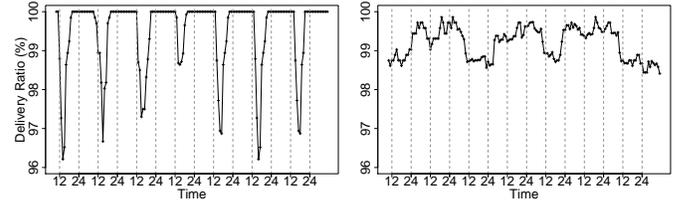


Fig. 10. Daily variations in *OPEN*, at high and low power. The sudden drop to zero corresponds to the 1-hour period in which, every two days, we downloaded the data from the flash of the nodes, and replaced their batteries.



(a) High power.



(b) Low power.

Fig. 11. Delivery ratio for CTP in *OPEN* (left) and *SPRUCE* (right).

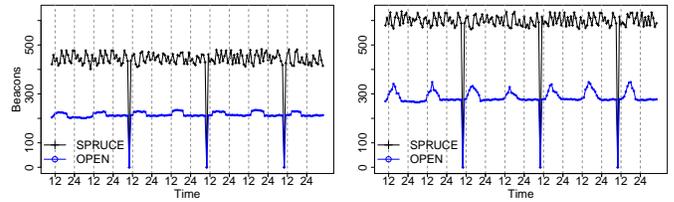


Fig. 12. CTP beacons in *OPEN* and *SPRUCE*, high (left) and low (right) power.

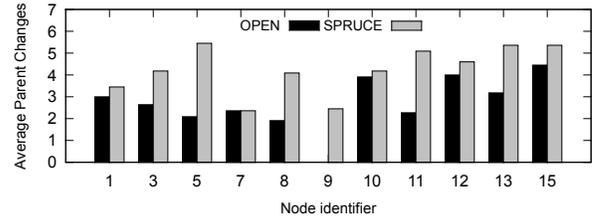


Fig. 13. Parent changes in *OPEN* and *SPRUCE*, at high power.

variations are much more frequent, and yield perfect reliability only for very short periods and only at high power. We argue that this behavior is induced by the environment. In *OPEN*, the absence of vegetation yields better and more stable links. Specifically, *i*) the link quality of the network in *SPRUCE* is lower than in *OPEN* (Figure 2) and, *ii*) in *SPRUCE* there are many bad links, while in *OPEN* there are a majority of good and perfect links (Figure 7(a) and 7(c)). On the other hand, in *SPRUCE* the negative effect induced by the vegetation is partially compensated by the fact that the latter “protects” links from environmental changes. In other words, links are worse on average, but subject to less abrupt variations.

Figure 12 shows instead the number of beacons as a function of time. Two observations can be made. First, this number is higher—almost double—for *SPRUCE* than in *OPEN*. Second, the variation over time in the former appears to be somewhat regular and independent of environmental variations (e.g., due to night and day), which are instead more visible in *OPEN*.

These observations can be explained with arguments similar to those for delivery rate. Indeed, CTP must frequently update the link quality information in *SPRUCE* simply because links are more brittle; we verified this hypothesis in our data sets,

where SPRUCE experiences many changes across link classes. This yields an increased energy expenditure for the control traffic required to maintain essentially the same (good) level of reliability of OPEN. In the latter, links are more stable; therefore, less updates are required, triggered to a greater degree by variations in the environment (e.g., day/night patterns).

These considerations are confirmed by looking at parent changes, shown in Figure 13 for high power only. Overall, the number of (average) parent changes we observed is exactly double—48 in SPRUCE vs. 24 in OPEN, due to the more challenging environment. However, nodes obviously experience different variations, due to the quality of their links. Daily patterns are less relevant here because the overall number of changes is smaller than in the case of beacons.

B. Contact detection for wildlife monitoring

We now report about experiments conducted with the hw/sw platform used in the wildlife monitoring application that, as mentioned in Section I, motivated our study. Unlike the staple data collection application we considered in Section VI-A, here nodes are animal-borne, and therefore mobile.

The goal is to detect *contacts* among animals, viz. the mobile nodes they carry. Nodes determine whether they are in proximity of others through periodic beaconing. Time is discretized into epochs, of length 60 s in our case, which determine the temporal resolution of contact detection. Each node sends a beacon at the epoch start and then listens for beacons from other nodes. A contact begins at a node upon receiving the first beacon, and ends when a beacon from the same node is not received within a given, user-defined time interval. In our case, the latter is equal to the epoch length; missing a single beacon causes the corresponding contact to be closed. We consider a contact detected when recorded by at least one of the two nodes involved.

Experimental setup. These tests were executed in the summer (July 11–August 4) of 2012. We deployed our application in the same sites described in Section IV, with the exception of OPEN which is the same as in Section VI-A. In the experiments we describe here, we wanted to assess the bias induced by the environment, therefore the application is run in a static topology, without the bias and complexity induced by mobility. We used the same cross topology described in Section IV. However, in this case we “stretched” it to obtain bigger distances, as this was required to test the application functionality, which was our primary goal when performing the experiments. This resulted in link distances up to 93 m, about 30% longer than in previous experiments.

The hardware we used is custom-made for our application, integrating a GPS unit and a GSM/GPRS modem. However, at its core is the CC2420 radio chip we used in the earlier experiments. The tests were performed only with high power. As in Section VI-A, we execute the real application: nodes are therefore not synchronized and collisions may occur.

Impact on the application layer. Figure 14 shows, for each environment, the number of detected and missed contacts. The

chart is built by discretizing time into 1-minute intervals (i.e., the epoch duration). We count, for each pair of nodes, whether in a given interval they are part of a contact (detected) or not (missed). The chart confirms the overall trends we reported in previous sections. The percentage of missed contact is lowest in OPEN and higher in the two forests, with SPRUCE better than BEECH. However, while the trend is the same (OPEN is better than SPRUCE, which is better than BEECH), the performance difference between environments is quite different w.r.t. the physical layer. In Figure 2(a), the *PDR* decrease in going from one environment to the other was 15-20%. Here, the difference in missed contacts between OPEN and SPRUCE is about 10%, while the one between SPRUCE and BEECH is about 40%.

The reason is that the two experiments measure different things in different setups. In Section V we analyzed the (aggregate) *PDR*, measuring directly packet transmission in a controlled scenario where collisions are absent. Here, we look at application data (with its own semantics, only partially related to packet transmission) in a much less controlled scenario. As a concrete example: missing a single packet causes a very small difference in the aggregate *PDR*, but it may determine a contact as closed in the application considered.

In other words, and similarly to what we already mentioned in Section VI-A about our data collection experiments, the observations about the physical layer in Section V in general cannot not be *directly* translated in quantitative terms to the application layer. However, there is clearly a relationship between the two, and knowledge of the quantitative tradeoffs in the former should inform design (if not prediction) of the latter, as we discuss in the next section.

VII. DISCUSSION AND OUTLOOK

Limitations of this study. The experiments we described in this paper required significant effort, as they had to be replicated in different outdoor sites, often in harsh conditions.

Nevertheless, more experiments are needed to confirm (or refute) the observations we presented. The findings we reported here are based on few relatively short experimental campaigns, and experiments were replicated across campaigns only for two of the three environments, and only during summer. Nevertheless, the close correspondence between the two campaigns, performed one year apart, appears very promising.

Other seasonal variations (i.e., in autumn and spring) and environmental conditions (e.g., rain or fog) could unveil additional trends and observations. Based on the promising results of this paper, we plan to extend the data set with these and other experiments. In this respect, our TRIDENT tool, supporting in-field tests without infrastructure, is an invaluable asset that greatly simplifies the experimental work.

Finally, other aspects could be considered. For instance, we collected plenty of data about RSSI and LQI, not presented here due to space limitations. These could be useful to

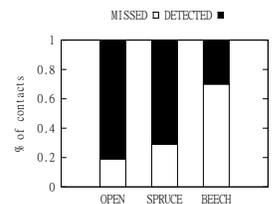


Fig. 14. The effect of environment on contact detection.

highlight long-term trends of the relationship between these metrics, commonly used as link estimators, and *PDR*. Also, we conducted experiments where the inter-packet interval (IPI) is fixed to 1 s. However, a smaller IPI is known to affect link properties [18]; we intend to investigate to what extent this phenomenon is affected by environmental conditions.

These limitations and improvements notwithstanding, to the best of our knowledge our study is the first one that looks at the problem of characterizing, from a quantitative standpoint, the overall behavior of a WSN in different environments. Next, we describe the research opportunities this opens.

Building on observations. This study is concerned with the *observation* of the impact of the environment on the physical layer. Nevertheless, end users (e.g., the biologists on our team) are concerned with the properties of the final application, e.g., in terms of data yield and lifetime. This paper clearly shows that there is a relation between the two. We showed, for instance, that the macro-trends across environments allow one to infer the relative trends in reliability and lifetime for CTP, and the likelihood of accurate contact detection.

These considerations can directly inform deployment decisions, and possibly application- or network-level strategies, from a *qualitative* point of view. For instance, several applications store locally the sensed data and only seldom report these data in bulk to the sink. If such an application is run in OPEN at low power, the considerations in Section V-D and Figure 11 suggest that performing the bulk transfers at night, when link quality is better, is likely to achieve better performance.

A more ambitious line of research is to use the observations we made in this work as the *quantitative* stepping stone enabling the *prediction* of application performance. In other words, the research question they inspire is: Can we define a methodology that, given the combination of environmental conditions relevant to the application, yields a good estimate of its performance, based on the principled execution and analysis of in-field tests? We already have the tools (e.g., TRIDENT) supporting the first step of the methodology. However, models are required to guide the use of tools (e.g., to minimize the duration of tests without sacrificing statistical relevance) and to link the results at the physical layer to the application. Answering this question would significantly improve the state of the art, by rendering the process of designing and deploying a WSN more repeatable, predictable, and less extemporaneous.

VIII. CONCLUSIONS

This paper presented an experimental study on the impact of environmental factors on outdoor WSN deployments. Motivated by a real-world application, we investigate the differences emerging when “immersing” the same WSN in different environments, characterized by varying degrees of vegetation, as well as seasonal and daily variations. We quantified the trends emerging at the physical layer, as well as their impact on the application layer. We argued that the observations we made in this paper can not only inform qualitatively the design and deployment of a WSN, but also inspire new methods to better estimate its performance in the target deployment environment.

REFERENCES

- [1] G. Anastasi, A. Falchi, A. Passarella, M. Conti, and E. Gregori. Performance measurements of motes sensor networks. In *Proc. of the 7th Int. Symp. on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, 2004.
- [2] N. Baccour, M. Ben Jamaa, D. do Rosario, A. Koubaa, H. Youssef, M. Alves, and L. B. Becker. A testbed for the evaluation of link quality estimators in wireless sensor networks. In *Proc. of the 8th Int. Conf. on Computer Systems and Applications (AICCSA)*, 2010.
- [3] N. Baccour, A. Koubaa, L. Mottola, M. Zuniga, H. Youssef, C. A. Boano, and M. Alves. Radio link quality estimation in wireless sensor networks: A survey. *ACM Trans. on Sensor Networks*, 8(4), 2012.
- [4] K. Bannister, G. Giorgetti, and E. K. S. Gupta. Wireless sensor networking for “hot” applications: Effects of temperature on signal strength, data collection and localization. In *Proc. of the 5th Workshop on Embedded Networked Sensors (HotEmNets)*, 2008.
- [5] C. A. Boano, J. Brown, N. Tsiftes, U. Roedig, and T. Voigt. The impact of temperature on outdoor industrial sensor network applications. *IEEE Trans. on Industrial Informatics*, 6(3), 2010.
- [6] F. Cagnacci, L. Boitani, R. A. Powell, and M. S. Boyce. Animal ecology meets GPS-based radiotelemetry: a perfect storm of opportunities and challenges. *Philosophical Trans. of The Royal Society Biological Sciences*, 365(1550), 2010.
- [7] M. Ceriotti, M. Chini, A. L. Murphy, G. P. Picco, F. Cagnacci, and B. Tolhurst. Motes in the jungle: Lessons learned from a short-term WSN deployment in the Ecuador cloud forest. In *Proc. of the 4th Int. Workshop on Real-World Wireless Sensor Networks Applications (REALWSN)*, 2010.
- [8] A. Cerpa, N. Busek, and D. Estrin. SCALE: A tool for simple connectivity assessment in lossy environments. Tech. rep. UCLA, 2003.
- [9] A. Cerpa, J. L. Wong, L. Kuang, M. Potkonjak, and D. Estrin. Statistical model of lossy links in wireless sensor networks. In *Proc. of the 4th Int. Conf. on Information Processing in Sensor Networks (IPSN)*, 2005.
- [10] M. Chini, M. Ceriotti, R. Marfievici, A. L. Murphy, and G. P. Picco. Demo: Trident, untethered observation of physical communication made to share. In *Proc. of the 9th Conf. on Embedded Networked Sensor Systems (SenSys)*, 2011.
- [11] F. Fabbri, M. Zuniga, D. Puccinelli, and P. Marrón. On the optimal blacklisting threshold for link selection in wireless sensor networks. In *Proc. of 9th Eur. Conf. on Wireless Sensor Networks (EWSN)*, 2012.
- [12] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks. Technical report, UCLA, 2002.
- [13] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis. Collection tree protocol. In *Proc. of 7th Conf. on Embedded Networked Sensor Systems (Sensys)*, 2009.
- [14] Y. Liu, Y. He, M. Li, J. Wang, K. Liu, and X. Li. Does Wireless Sensor Network Scale? A Measurement Study on GreenOrbs. In *Proc. of 30th Int. Conf. on Computer Communications (INFOCOM)*, 2011.
- [15] L. Mottola, G. P. Picco, M. Ceriotti, S. Guna, and A. L. Murphy. Not all wireless sensor networks are created equal: A comparative study on tunnels. *ACM Trans. on Sensor Networks*, 7(2), 2010.
- [16] R. Musaloiu, C.-J. M. Liang, and A. Terzis. Koala: Ultra-low power data retrieval in wireless sensor networks. In *Proc. of the 7th Int. Conf. on Information Processing in Sensor Networks (IPSN)*, 2008.
- [17] D. Son, B. Krishnamachari, and J. Heidemann. Experimental analysis of concurrent packet transmissions in wireless sensor networks. In *Proc. of the 4th Conf. on Embedded Networked Sensor Systems (SenSys)*, 2006.
- [18] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis. An empirical study of low-power wireless. *ACM Trans. on Sensor Networks*, 6(2), 2010.
- [19] K. Srinivasan, M. A. Kazandjieva, M. Jain, E. Kim, and P. Levis. SWAT: enabling wireless network measurements. In *Proc. of the 6th Int. Conf. on Embedded Networked Sensor Systems (SenSys)*, 2008.
- [20] J. Thelen and D. Goense. Radio wave propagation in potato fields. In *Proc. of the 1st Workshop on Wireless Network Measurements*, 2005.
- [21] J. Zhao and R. Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *Proc. of the 1st Int. Conf. on Embedded Networked Sensor Systems (SenSys)*, 2003.
- [22] M. Zuniga and B. Krishnamachari. An analysis of unreliability and asymmetry in low-power wireless links. *ACM Trans. on Sensor Networks*, 3(2), 2007.