

Motes in the Jungle: Lessons Learned from a Short-term WSN Deployment in the Ecuador Cloud Forest

Matteo Ceriotti¹, Matteo Chini²,
Amy L. Murphy¹, Gian Pietro Picco², Francesca Cagnacci³, and Bryony Tolhurst⁴

¹ Fondazione Bruno Kessler—IRST, Trento, Italy {ceriotti, murphy}@fbk.eu

² Dip. di Ingegneria e Scienza dell'Informazione (DISI), Univ. of Trento, Italy
matteo.chini@gmail.com, gianpietro.picco@unitn.it

³ Edmund Mach Foundation—IASMA, S. Michele all'Adige, Italy
francesca.cagnacci@iasma.it

⁴ Biology Division, School of Pharmacy and Biomolecular Sciences, Univ. of Brighton, UK
bryonytolhurst@live.co.uk

Abstract. We study the characteristics of the communication links of a wireless sensor network in a tropical cloud forest in Ecuador, in the context of a wildlife monitoring application. Thick vegetation and high humidity are in principle a challenge for the IEEE 802.15.4 radio we employed. We performed experiments with stationary-only nodes as well as in combination with mobile ones. Due to logistics, all the experiments were performed in isolation by the biologists on our team. In addition to discussing the characteristics of links in this previously unstudied environment, we also discuss the lessons we learned from operating under peculiar constraints in a peculiar deployment scenario.

1 Introduction

Wireless sensor networks (WSNs) are applied in many scenarios, each with unique characteristics in terms of connectivity. Assessing the specifics of a target environment is usually complex, and often entails a preliminary pilot deployment.

Application context and motivation. In this paper we report about such a pilot deployment, which took place in the cloud forest of the North-Western slopes of Ecuadorian Andes during March 29–April 3, 2010, and whose details are provided in Section 2.

The work described here is part of a larger research effort targeting the monitoring of biodiversity in community-based primary cloud forest reserves in this Andean region. Indeed, this area is at the confluence of two of the world's hottest biological hotspots: the Chocó-Darién Western Ecuadorian and the Tropical Andes. Available checklists of vertebrates likely miss most reptile and mammal species, including medium-to-large ones. The knowledge about these species' use of space and community interactions is essential to ascertain their susceptibility to environmental changes and guide conservation measures. Available information is extremely sparse and based on discontinuous observations and occasional surveys. Direct observation of animals is not a robust method, due to the very dense vegetation, while traditional indirect methods, such as capture-mark-recapture or radio-tracking are extremely effort-demanding as these areas are secluded. Recent advancements in wildlife studies, e.g., the use of GPS devices, are expensive and therefore applicable to a small number of species and sample size. WSNs

provide a new, exciting option in such challenging environmental conditions, especially for long-term monitoring. Advantages include the need for only a single capture (to fit the node) and the possibility to study a large sample thanks to the relatively low equipment and deployment cost. However, an essential step in seizing this opportunity is the evaluation of the node performance in the target environment.

The envisioned WSN application will encompass nodes permanently deployed in the environment at known locations as well as attached with collars to the animals themselves. We intend to use motes functionally equivalent to Moteiv’s TMote Sky [3], arguably the most popular platform today. However, the 2.4 GHz band used by the CC2420 radio chip on these motes is known to be highly sensitive to foliage and water—essential ingredients of a cloud forest. Therefore, the primary motivation behind the study described here was to assess the connectivity characteristics of the target environment to determine the feasibility of our WSN architecture and guide its design.

Related work. A few real-world deployments focus on forests [5], but with characteristics different from ours. Despite the importance of understanding the connectivity of the environment targeted by a WSN, this information is rarely reported in the literature. Instead, the problem is usually tackled with studies targeting either static [4] or mobile [1] scenarios. All the reported works, however, leverage the possibility to progressively refine the investigation based on the findings. Our need to define a priori the entire experimentation pushed us towards a more general methodology, something still not available in the literature. To design our study we leveraged our prior expertise in comparing the network characteristics of a tunnel against the vineyard environment [2]. However, the differences in the application scenario, involving mobile nodes, and the inability to access the experiment site demanded a significant revision of our techniques.

Challenges. The deployment itself presented non-trivial logistical difficulties due to the geographical distance and the harshness of our target environment. Things were further complicated by the fact that the WSN experiments were “piggybacked” on the biologist’s trip to Ecuador for other research purposes.

As a consequence, we faced rather unusual requirements. In the literature, similar experiments are typically run by the WSN developers, often in rather controlled environments. Instead, in our case the experiments had to be run by the biologists, and *in isolation*. Remote WSN configuration was not an option, due to the absence of data connectivity from the experiment location—the jungle. Similarly, a multi-phase deployment, where the output of one experiment guides the setup of the next, was also not an option due to the distance between the experiment location and the closest Internet access, and to the duration of the experiments. The latter was limited by the biologist’s already-established trip schedule, further reduced by the inevitable lost baggage.

Simply put, this meant that our hw/sw WSN setup had to work out of the box for the entire duration of the experimental campaign, and had to be simple enough to be operated by someone without expertise with this technology.

Contributions and findings. The details about our cloud forest experiments are provided in Section 3. The main contributions of this paper are the following:

1. *Low-power wireless in the jungle environment.* In Section 4 we analyze the gathered data. The depth of the analysis is somewhat limited by the aforementioned logistic problems, as we did not have a second chance to investigate the source of unexpected behaviors. However, we are not aware of other studies investigating

low-power wireless communication in an environment similar to ours and therefore, even with these limitations, we believe our study can be of value for the research community. Moreover, some of our findings are somewhat surprising. For instance, we expected links to be rather short and unreliable, due to foliage, water, and humidity. Instead, our data show that 30-meter links are common, and in some cases reliable communication occurs up to 40 m.

2. *Mobile nodes as a connectivity exploration tool.* The inclusion of experiments with mobile nodes was initially motivated by the animal-borne nodes in our envisioned application. We expected to draw the bulk of our considerations from stationary-only experiments. Instead, mobile nodes played a much more relevant role in our study. On one hand, the stationary-only experiments did not deliver the amount of data we expected. The connectivity patterns were not known in advance, and a multi-phase deployment was not an option, as already discussed. Mobile experiments provided a data set complementing the stationary ones. On the other hand, with hindsight, the use of mobile nodes is an effective way to explore connectivity, regardless of mobility requirements. Intuitively, a broadcasting node moving through a single, well-designed path yields a wealth of information, more varied and fine-grained w.r.t. stationary-only experiments, even considering the interference introduced by the person executing the experiments. This enables a more precise “connectivity map” of the environment, that can be used for instance to guide node placement. We believe the use of mobile nodes can become an essential element of studies aimed at characterizing connectivity in WSN environments.
3. *When WSN developers are not in charge.* Our experiments were run by someone other than the WSN developers because of opportunity. There may be other reasons, e.g., the necessity to require authorizations or safety concerns related to the target deployment area. In any case, for WSN to become truly pervasive, end-users must be empowered with the ability to deploy their own system. The lessons we learned, distilled in Section 5, can be regarded as a contribution towards this goal.

2 Deployment Scenario

Location. The community-based reserve of Junin, in the Intag region of the Imbabura province in Ecuador ($0^{\circ}16'19.09''\text{N}$; $78^{\circ}39'28.92''\text{W}$) is between 1,200 and 2,800 m above sea level of the North-Western slopes of the Ecuadorian Andes. Significant portions of these mountain areas are primary cloud tropical forests, almost permanently cloudy and foggy. According to the United Nation’s World Conservation Center, cloud forests comprise only 2.5% of the world’s tropical forests, and approximately 25% are found in the Andean region. Therefore, they are considered at the top of the list of threatened ecosystems. The climate is tropical, and the flora and fauna incredibly rich, with about 400 species of birds and 50 known mammal species (including 20 carnivores), many probably still unchecked or even unknown. The small human community of about 50 people is 20 km from the closest village, and a 7 hour dirt-road drive from the closest town. The vegetation is made by relatively scattered mature trees, constituting the canopy, and a dense undergrowth of shrubs and epiphytes. During the rainy season (November-May), when we ran our experiments, it rains every day for nearly the entire day.



Fig. 1. Packaging.



Fig. 2. In the jungle with mobile nodes.

WSN Equipment. Our experiments used 18 TMote Sky nodes, equipped with the Chip-Con 2420 IEEE 802.15.4-compliant, 2.4 GHz radio and on-board inverted-F micro-strip omni-directional antenna. The choice of this popular platform is motivated both by our intended use of a similar platform in our own wildlife application, and to enable comparison with similar experiments in different environments reported in the literature. Alternate hardware would significantly modify the results, e.g., an external antenna would likely dramatically increase the observed connectivity. Moreover, these nodes are provided with an external flash memory, enabling storage of the experiment data.

As stationary nodes were intended to be attached to trees in a very humid environment, under heavy rain, we used IP65 water-proof boxes with a transparent cover, enabling the sampling of the light as requested by the biologists. Inside each box we glued a USB female connector to easily anchor and replace the node as needed. Each box also contained a battery holder with two size D batteries and desiccant bags to protect the node against humidity. The packaging is shown in Figure 1 in the same orientation as it was attached to the trees. In contrast, the mobile node was simply a TMote Sky powered by 2 AA batteries, wrapped in a plastic bag.

3 Experiment Design

The WSN was composed of 8 nodes, placed in a cross configuration, as shown in Figure 3(a). The placement was determined as part of the stationary experiments, described next. Node 0 served as the experiment coordinator, broadcasting a message indicating the start time and configuration of each experiment. All communication took place on channel 18. Since no computer was available in-field, we used the nodes' LEDs to visualize the node functionality. For example, toggling the yellow LED indicated message transmission, while toggling together the other two LEDs indicated message reception. At node boot time, a visual code for the battery voltage was shown to advise for battery




Fig. 3. Deployment of stationary nodes; each color corresponds to about 1 m difference.

replacement in case of values below 2.7 V, the minimum required to write to the flash memory. To start an experiment, the biologist pressed the user button.

The software was built on TinyOS and without any MAC protocol, given our goal of characterizing physical connectivity. Packet collision was avoided by an appropriate transmission schedule sent at the beginning of each experiment by node 0. For each experiment, and for each link $i \rightarrow j$, we recorded in the flash the following metrics:

- *Packet Delivery Ratio* ($PDR_{i \rightarrow j}$), the number of packets received at node j over the total number of packets sent by node i ;
- *Received Signal Strength Indicator* ($RSSI_{i \rightarrow j}$), the signal strength of the packets transmitted by i , as observed by the radio of j ;
- *Link Quality Indicator* ($LQI_{i \rightarrow j}$), the correlation index between the symbol received at j , sent by i , and the one to which it is mapped after radio soft decoding.

3.1 Preliminary Tests

Goals. Given the lack of reported experiences in scenarios similar to ours, the primary goal of these tests was to determine the communication range, to properly place nodes in the next experiments. These experiments also investigated different power transmission levels as well as the impact of direct tree obstruction.

Implementation. The experiments exploited only node 0 and 3 in Figure 3(a). We implemented two experiments, one to determine the range of communication, and the other to investigate the effect of signal power and tree obstruction. In the former, each node sent 600 messages with an inter message interval (IMI) of 2 s. All messages were sent with -1 dBm transmission power. The LED visual feedback was used to guide the identification of the maximal communication range. In the latter experiment, each node sent a sequence of 3000 messages with a 2 s IMI, interleaving sending between the involved nodes. These messages are logically divided into 5 tests of 600 messages each, 3 at -1 dBm, commonly used in WSN deployments, and 2 at -8 dBm, to investigate the effect of reduced power. For each 600-message set we stored the aggregated average RSSI (\overline{RSSI}), average LQI (\overline{LQI}) and PDR values over all received messages.

Deployment. In all experiments node 0 was attached to a tree at 1 m height, while node 3 was placed on a chair. In the first experiment, the two nodes were in line of sight (LoS) and the biologist gradually moved the chair away from the tree while monitoring the LEDs for determining a safe communication range, which she established at 28 m. The second experiment with different power levels was run a first time with nodes in line of sight, and then again with node 0 directly behind the tree, creating a link obstruction.

3.2 Tests with Stationary Nodes

Goals. The purpose of these tests was to investigate connectivity among nodes at different distances, over a long time interval, and at different node heights.

Implementation. These experiments used the nodes as in Figure 3(a) and, as in the preliminary tests, relied on node 0 for disseminating the start time and transmission schedule. In each experiment, each node sent 215 messages with an IMI of 8 s, resulting in an interval of 1 s between nodes adjacent in the transmission schedule. The experiments

were batched and ran for an entire day, interleaving 23 experiments at -1 dBm with 22 experiments at -8 dBm. Before this batch, a 1-hour *setup* experiment (with LEDs enabled) was performed, to verify connectivity and thus node placement. At the end, each node computed and stored the overall PDR , \overline{RSSI} , and \overline{LQI} w.r.t. all other nodes.

Deployment. Node 0 and 3 were left in place after the preliminary tests. During the *setup* experiment, all the others were moved one by one away from node 0 in small steps. Based on high-level instructions, the LEDs blinking, and the communication range of 28 m determined in the preliminary tests, the biologists determined the final placement shown in Figure 3(a), yielding the set of distances covered as shown in Figure 3(b). The experiments were executed twice for a total of 2 days.

Our original idea was to deploy the nodes in a flat area, placing them first at ground level, then at 1 m from the ground, and finally at various, possibly higher heights. The rationale was to determine node placement in the least favorable connectivity conditions, close to the ground. Unfortunately, due to the delayed arrival on site (caused by lost luggage), the biologists decided to eliminate the first experiment. Moreover, due to the available terrain, highly irregular and on a sort of hill as shown in Figure 3(a), the second and third deployments were reversed. Therefore, the deployment was setup in the connectivity conditions *most* favorable, which affected the subsequent experiments. Indeed, undergrowth interfered significantly during the second test, making its results unusable. Also, node 2 failed to start some tests and its data has been excluded.

3.3 Tests with Stationary and Mobile Nodes

Goals. These experiments were initially motivated by our wildlife application, combining fixed and animal-borne nodes. When interpreting the results, however, we realized the importance of these tests in enabling exploration of connectivity at many more distances w.r.t. the static deployment, yielding more spatial continuity to data points.

Implementation. In these experiments, node 0 was carried by the biologist, who moved throughout the deployment area. Stationary nodes only listened, while node 0 broadcast messages at -1 dBm for 15 min, with an IMI of 500 ms, yielding 1,800 messages per experiment. Unlike stationary experiments, which recorded only one aggregate value for each link, in the mobile tests statistics about each individual message were recorded. This allowed us to treat each message separately, by considering the distance between the mobile node and each stationary node at the moment it was sent. Offline data correlation across nodes was enabled by timestamping the message at the sender, and saving this along with the RSSI and LQI values at the receiver. During experiments the biologist moved freely, her path recorded by a video camera carried by a second team member (Figure 2), allowing us to visualize the movements and correlate the timings.

Deployment. The placement of stationary nodes was the same as in Section 3.2, but the nodes were physically replaced as their (pre-loaded) software was different. The nodes were placed at 1 m from the ground. The mobile node was either held in the biologist hands (as in Figure 2) with the antenna parallel to her shoulders and the board facing the sky or carried chest height inside a pouch, unfortunately with undefined orientation. First, the biologist stood near a stationary node (node 2) and made simple movements of approximately 1 m amplitude along the horizontal plane at the node height and along the tree, approaching the node from four directions—front, back, right, and left. Then, the

Link	TX power	PDR		\overline{RSSI}		\overline{LQI}	
		LoS	Tree	LoS	Tree	LoS	Tree
0 \rightarrow 3	-1 dBm	86.7%	79.5%	-87 dBm	-91 dBm	99	90
3 \rightarrow 0	-1 dBm	84.4%	69.7%	-88 dBm	-92 dBm	98	88
0 \rightarrow 3	-8 dBm	24.2%	1.3%	-92 dBm	-93 dBm	80	77
3 \rightarrow 0	-8 dBm	11.8%	0.5%	-92 dBm	-94 dBm	77	75

Table 1. Results from the preliminary tests.

biologist moved back and forth between node 1 and 3, then between 2 and 5. Although these experiments focused on movement between a subset of the available nodes, all nodes in the network recorded message reception, thus we gathered a large amount of data. Finally, the biologist composed a path visiting all stationary nodes. Each path was repeated 4 times. In total, these experiments produced 116,448 data points. We excluded the data collected by node 7 as we verified that its short-range reception was abnormal.

4 A Mote's Life In the Jungle

4.1 Preliminary Tests

The results of the tests on transmission power and tree influence are shown in Table 1. As discussed in Section 3.1, these involved only node 0 and 3. At -1 dBm, both PDR and \overline{LQI} are high. This is expected as these results are at the distance of 28 m the biologist chose as the border of good connectivity. Interestingly, our initial guess for a safe communication distance was much lower, around 10-15 m, given the presence of thick vegetation and high humidity. \overline{RSSI} is low but, given the absence of radio interference in the forest, it does not significantly affect PDR . The presence of a tree right in front of a node may cause link asymmetries. With nodes in line of sight, the PDR difference between the two link directions is only 2%, but with the tree in between this increases to 10%, indicating a weaker link when communication originates near the tree. \overline{RSSI} and \overline{LQI} do not show marked asymmetries, although they decrease when the tree obstructs the link. With lower transmission power, PDR is non-negligible but more heavily influenced by the tree. The low \overline{LQI} is consistent with the next experiments showing that 28 m is well outside the good-connectivity range at -8 dBm.

4.2 Tests with Stationary Nodes

Long-Distance, High-Quality Links. We expected the dense jungle foliage to significantly limit communication. Instead, Figure 4(a) shows that communication is almost perfect up to 20 m, although the high PDR at 19.8 m occurs with a relatively low signal strength (Figure 4(b)). Further, although the 38 m link falls well beyond the region with




Fig. 4. Average and standard deviation of the results from stationary tests with power -1 dBm.




Fig. 5. *PDR* over time with power -1 dBm from stationary tests.




Fig. 6. Results over time with power -8 dBm from stationary tests.

perfect communication, analysis over time (Figure 5) shows that this link was also perfect for more than half of the experiment duration. While this is clearly an anomaly of the setup, it clearly demonstrates that connectivity in the jungle is much different than expected. At -8 dBm, the area with perfect links is only slightly reduced to 14 m.

Fluctuations and Asymmetries of Mid-Range Links. Figure 4(a) and 4(b) show that links with mid-range distances of 20–40 m have highly-variable quality and low RSSI. The *PDR* large standard deviation is best viewed over time in Figure 5, where each point describes the result of one 30-min experiment for a given link. From the detail on the right-hand side of the figure, one can see that the variability is unpredictable. For example, around hour 15 some links improve while others decline. Further, some links such as (3, 0) show transient asymmetries. Weather could be the culprit, and indeed it rained during the majority of these tests. Although one would expect a global decay of link quality, it is possible that humidity, rain, and pools of collected water affect communication in local, unpredictable ways, although we do not have direct observations confirming this. In any case, mid-range links clearly cannot guarantee connectivity, but they can certainly be exploited transiently by adaptive routing algorithms.

Long-Range Interference with Reduced Power. At -8 dBm, links outside the perfect communication range disappear for long periods of time (Figure 6). While these links are basically unusable, they can cause long-range interference. For example, Figure 6(b) shows messages received with very low RSSI even at 40 m. Although these distant transmissions rarely succeed, they could easily disrupt overlapping shorter-range ones.

4.3 Tests with Stationary and Mobile Nodes

“Omnidirectional” Antenna. Figure 7 shows the effect of a node approaching a second one fixed to a tree, as described in Section 3.3. Based on the biologist’s 1-meter




Fig. 7. Node 0 approaching node 2, attached to a tree, from different directions.

horizontal movements, the different shapes of the *Front*, *Left*, and *Back* curves clearly show the well-known fact that the used antenna is not perfectly isotropic. Interestingly, the flat tops in *Right* do not correspond to a movement pause, rather to the “saturation” of RSSI for very short links. Tree obstruction is clearly evident in the *Back* curve.

Influence of Body, Tree, and Ground. In Figure 8 the biologist, holding the mobile node in front of her chest, looped four times around nodes 1 and 3. We decomposed the data trace to distinguish the possible obstructions. For example, when walking from 1 to 3, the tree obstructed communication received at 3 (Figure 8(b)), and the body obstructed receptions at 1 (Figure 8(c)). As a reference, we chose the line-of-sight case: reception at 1 when walking from 3 to 1 (Figure 8(a)). The same experiment was run with the mobile node held a few centimeters from the ground (Figure 8(d)).




Fig. 8. Effect of tree, body, and ground on communication. The line in the RSSI plots shows the delta in percent w.r.t. the line-of-sight shown in (a).

Trees induce a reduction up to 20% on \overline{RSSI} in short links (< 20 m), while longer links are not affected. The body also reduces \overline{RSSI} in short links, but more significantly, up to 40%. Moreover, the body reduces the maximum communication range by 10 m, as denoted in Figure 8(c) by a nearly-zero PDR beyond 30 m. As expected, the simultaneous obstruction of tree and body, not shown for space reasons, yields a combination of previous results: a shorter communication range and \overline{RSSI} reductions up to 60%. This bears an important implication for our wildlife application, where we need to estimate the distance between animals upon contact: RSSI-based distance approximation schemes may have a significant error, induced by trees, the body of animals, and the direction the animal approaches the tree, as discussed previously.

Placing the sender near the ground produces a different combination of effects. Specifically, the line-of-sight communication range is much shorter than in Figure 8(a), but the \overline{RSSI} is affected by at most 20%. As this scenario is the closest to our target deployment with tagged animals, it warrants additional study.

4.4 An Evaluation of Mobile Nodes as Connectivity Probes

We take a step back from the data analysis to consider our data collection methodology, specifically, comparing the results of stationary test against those with mobile ones.

Aggregated Mobile Tests vs. Stationary Tests. Thus far we have looked only at excerpts of the mobile traces, extracting cases with specific characteristics. Here, we aggregate all data points collected over all node movements, with the results shown in Figure 9(a)–9(c). To plot PDR , we calculate the distance between the mobile and each stationary node, then plot the number of messages received over those sent at each distance. \overline{RSSI} and \overline{LQI} are instead shown as the average and standard deviation over all




Fig. 9. Aggregated results over all 11 mobile experiments. In (d) and (e), the difference in PDR for the links longer than 38 m is outside of the chart range.

the messages received along links of a specific length. We then compare these data to those collected in the stationary tests of Figure 4, by plotting the percentage difference in Figure 9(d), only for the points studied in the stationary scenario.

In the mobile scenario, the reduction of \overline{RSSI} on short links (< 15 m) is likely attributable to body interference as observed in Figure 8(c). From the PDR comparison in Figure 9(a), we note that at all distances, the mobile scenario produces *worse* results, meaning that the PDR at a given distance is lower in the mobile scenario than in the stationary. To understand the implications, consider that we intend to use the results of this study to plan a future deployment. If we base this deployment only on the results of the mobile study, all stationary nodes in our future deployment would certainly be connected. Instead, if we base our fixed node placement on the stationary results, we would erroneously expect to communicate with mobile nodes carried by animals at the same distance. In other words, the mobile case underestimates the communication potential of stationary nodes while the stationary overestimates communication to mobile nodes.

Interestingly, Figure 9(c) shows better quality links in the mobile scenario. While this is opposite from the observations of PDR , the stationary experiments showed that LQI varied significantly throughout the day. Instead, the mobile experiments were concentrated in less time, and may have taken place in favorable connectivity conditions.

Figure 9(e) accounts only for the data recorded in conditions similar to those of the stationary only tests, i.e. removing the body shielding and using the data from Figures 8(a) and 8(b), namely LoS and tree-only obstruction. For short links (< 20 m), values are in agreement while longer links are hampered by interference from the ground and dense low-level foliage in the mobile scenario. In the stationary tests, nodes were always within LoS, therefore the undergrowth had minimal effect.

Statistical Relevance of Mobile Tests. The experiments run with a mobile node made it possible to explore the physical space in a continuous fashion, spreading the collected data points over more distances w.r.t. stationary-only tests. To understand the effectiveness of this approach, Figure 10 compares the average number of messages received in 1 hour for each distance covered in the mobile case, to the number of messages that the stationary experiment would receive with the same IMI as the mobile nodes, i.e. 500 ms. Recall that to avoid collisions, our stationary experiments used a 1 s IMI. The distribution of the tested distances is naturally biased by the executed movements. Nonetheless, even without a guided motion plan, all distances less than 40 m have been tested by at least 400 messages, i.e., 25% of the messages sent by the stationary tests for each link. The ability of the mobile node to cover so many distances clearly motivates its use as a probe to characterize connectivity.




Fig. 10. Number of messages sent per hour along links of a given distance by stationary and mobile tests.

5 Lessons Learned and Future Work

Our experiments were run in a challenging scenario by biologists without WSN expertise, with limited equipment, and in isolation. We have never faced this combination in our previous real world deployments, and we learned interesting lessons.

Mobile Nodes: Application Insights or Connectivity Probes? It was the biologists who requested experiments with mobile nodes, to concretely understand what WSNs could offer them. Nevertheless, we learned that the use of mobile nodes, despite the inherent imprecision, is useful for characterizing an unknown environment and guiding the actual deployment. Further work is needed to explore the opportunities of this technique and understand its limitations, e.g., the difficulty to capture long-term variations.

The Role of LEDs. In our study, the node output had to be simple yet informative enough to guide the biologists. Our solution, based on giving a visual clue only about send/receive operations, contributed to the creation of very long links between stationary nodes which in turn contributed to the failure of the second set of stationary experiments, as mentioned in Section 4.2. A visual representation of the RSSI values (e.g., represented by a “histogram” using the three LEDs), would have led to shorter links, which would have produced meaningful data even in the second set of experiments.

Testing Blindly. Our experimental campaign involved many decisions taken blindly. We did not have an understanding of the environment based on previous studies. We did not have a well-defined methodology for performing this kind of experiments, and none yet exists in the WSN field. Finally, we could not modify experiments based on intermediate results. We partially reduced the unknowns by breaking down experiments into phases with well-defined outputs. Examples are the preliminary tests (Section 3.1) and the 1-hour *setup* phase preceding the stationary tests (Section 3.2). These enabled the biologists to take informed decisions autonomously, partially obviating the absence of WSN experts in-field. Nevertheless, this did not avoid incorrect decisions, and could not provide answers for unanticipated questions (e.g., the cause of high time variance of links). How much can we reconcile the autonomous execution of experiments and the depth of the resulting analysis? To what extent can we automate the process? These are interesting research questions and the subject of our ongoing work.

Acknowledgments. This work has been partially supported by the EU Cooperating Objects Network of Excellence (CONET: FP7-2007-2-224053) and by Martin Stanley (Holly Hill Trust). We are also indebted to the Junin community, in particular Rosario Piedra, Victor Hugo Ramirez, Olga Cultid, for logistical support, and exceptional hospitality.

References

1. Miluzzo, E., Zheng, X., Fodor, K., Campbell, A.T.: Radio characterization of 802.15.4 and its impact on the design of mobile sensor networks. In: Proc. of the 5th European Conf. on Wireless Sensor Networks (EWSN) (2008)
2. Mottola, L., Picco, G., Ceriotti, M., Guna, S., Murphy, A.: Not All Wireless Sensor Networks Are Created Equal: A Comparative Study on Tunnels. ACM Trans. on Sensor Networks (TOSN) 7(2) (2010)
3. Polastre, J., Szewczyk, R., Culler, D.: Telos: Enabling ultra-low power wireless research. In: Proc. of the 5th Int. Conf. on Information Processing in Sensor Networks (IPSN) (2005)
4. Srinivasan, K., Dutta, P., Tavakoli, A., Levis, P.: An empirical study of low-power wireless. ACM Trans. on Sensor Networks (TOSN) 6(2), 1–49 (2010)
5. Tolle, G., Polastre, J., Szewczyk, R., Culler, D., Turner, N., Tu, K., Burgess, S., Dawson, T., Buonadonna, P., Gay, D., Hong, W.: A macroscope in the redwoods. In: Proc. of the 3rd Int. Conf. on Embedded Networked Sensor Systems (SenSys) (2005)