Experimental Analysis and Characterization of a Wireless Sensor Network Environment

Bogdan Pavkovic
Grenoble Informatics Laboratory (LIG)
University of Grenoble
pavkovic@imag.fr

Fabrice Theoleyre
CNRS, LSIIT
University of Strasbourg
theoleyre@unistra.fr

Dominique Barthel
Orange Labs, Meylan
dominique.barthel@orange-ftgroup.com

Andrzej Duda
Grenoble Informatics Laboratory (LIG)
University of Grenoble
duda@imag.fr

ABSTRACT

Existing testbeds, even though rare and specialized, are often not used to their full potential. Collected data during an experimentation is usually used to evaluate some specific tested aspect. To further benefit from the knowledge gathered on a testbed and to obtain the insight into the WSN environment itself, we propose a thorough statistical analysis. Some of our analysis include radio link characterization, its correlation with environmental parameters as well as an insight into network dynamics from the point of view of a node and a link. We also discuss how testbeds should be designed or improved to provide more detailed information necessary for an advanced analysis.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: [Network Architecture and Design - Wireless Communication]

General Terms

Experimentation, Algorithms, Performance

Keywords

wireless sensor networks, testbeds, characterization, statistical analysis

1. INTRODUCTION

Wireless Sensor Networks have attracted a lot of attention for a few years. Most of the research effort has focused on the two main issues: routing [1] and power-energy savings [2]. However, the research community has become aware that models of wireless multihop networks are too simplistic and lead to misleading conclusions. In particular, different simulators have been proven to provide different results [3]. The radio model especially has a strong impact on performance [4].

To advance the evaluation of various protocols, we can setup an *ad hoc* testbed to compare the simulation results with measurements gathered on the testbed. However, the measurements usually concern only a limited number of tested aspects and setting up operational testbeds requires large human efforts.

Developing new testbeds to accelerate prototyping and to make their evaluation easier has become an interesting research objective. Orbitlab [5] proposed for example to deploy a testbed in a dynamically reconfigurable grid. Other authors provide guidelines to design feasible protocols: they advocate for the principle of *the simplest is the best* [6]. However, these commendable efforts usually do not provide generic results to the networking research community. For instance, they do not consider many important aspects such as: What are the characteristics of the WSN radio topology? What is the reliability of a WSN? Are the properties stable or do they exhibit some variability or periodicity? We propose here to address one part of these fundamental concerns.

In the past, statistical analysis has been applied to traffic analysis [7] or anomaly detection [8] to extract some correlations and salient features. We propose in this paper to use this mathematical approach to characterize the performances and properties of a WSN. Our analysis includes in particular:

- characterization of radio links in a WSN: their reliability and the correlation between their properties;
- analysis of the network dynamics: how does a WSN change in time?
- how can we predict the quality of a radio link with a local and simple measure?
- how can we discard measurement errors?

2. TESTBED DESCRIPTION

We used a testbed originally designed for validating a routing protocol [9]. It was composed of 36 Coronis nodes

PE-WASUN'10, October 17–18, 2010, Bodrum, Turkey.

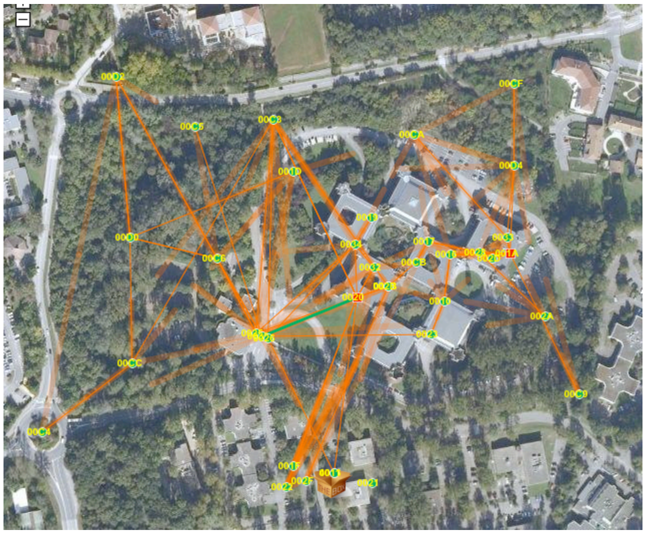

Figure 1: Deployed topology in an urban environment

implementing the Wavenis technology [10]: a fast frequency hopping is implemented to be robust to narrow band interference. Nodes operate in the 868 MHz license-free band, emitting at 25 mW with maximum transmission rate of 19200 bps. The MAC layer follows a CSMA-CA approach for medium access contention. Besides, two nodes acted as sinks with a direct connection to the Internet and a database for storing received packets. Nodes were deployed over the area of the technical park of Orange Labs in Meylan, France, both indoor and outdoor. Their location is diversified enough (e.g. walls, barrier, trees, ceiling) so that a wide range of situations is observed. We analyzed the measurements of 18 days of operation. Figure 1 presents the deployed topology in the urban environment.

The testbed was mainly used to validate a routing protocol based on virtual coordinates: each node maintains a metric related to its virtual distance to the sink [11]. The next hop is chosen as the neighbor that is virtually the closest to the sink.

Nodes perform a neighborhood discovery every 13 minutes and maintain a proactive neighborhood table including the virtual distance and RSSI of each neighbor.

Each node generates a new data packet every 17 minutes. This packet is transmitted in anycast: any sink can be used to reach the wired part of the network. In order to select the next hop (the node that has the lowest virtual distance), a node only has to walk its neighborhood table.

The routed packets, aside from the control fields (source and destination ID, sequence number, etc.), contain debug information consisting of complete neighborhood tables (the neighbor ID and the received RSSI value: 32 possible levels between -108 dBm and -60 dBm in 1.5 dBm increments) and the application payload consisting of measurements of the temperature, humidity, and light sensors at the instant just before sending the packet. Packets successfully received at sink nodes were labeled with a timestamp and stored in a database. Table 1 sums-up the important testbed information.

3. METHODOLOGY

Environment type	Urban
Node position	Indoor & outdoor
Sensor type	Coronis Wavenis
Number of nodes (sinks)	36(2)
Duration of the experiment	18 days
Neigh. discovery period	13 min.
Data packet generation period	17 min.

Table 1: Testbed parameters

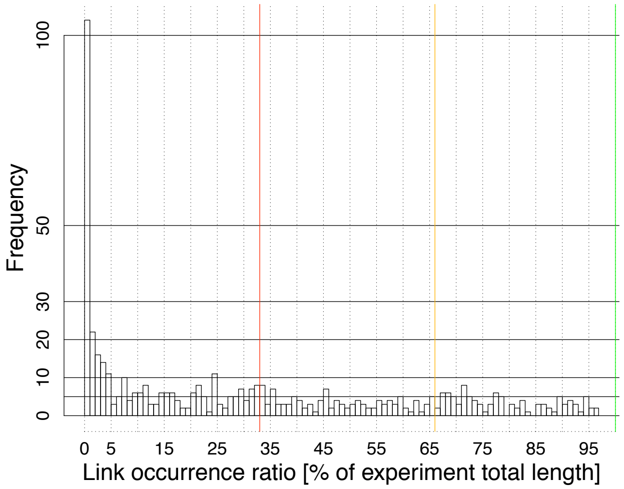

Figure 2: Distribution of link occurrence ratio

3.1 Database description

To allow meaningful interpretation and easy use of different types of measured values contained in the received routing packets, the database is divided into few tables:

- the node ID and its geographical position (known prior to deployment) to obtain the geographical topology. We can compare it to the radio topology;
- neighborhood information (neighbor ID and a RSSI value). We can observe in particular duration and quality of the links;
- sensor measurements (e.g. humidity, temperature).

On the average, each node sent 1,500 data packets (maximum sample size) to sinks, where just the ones successfully arrived were saved in the database. To perform an accurate statistical analysis, we need to discard received data samples with insufficient cardinality. Thus, we have removed all the data samples that count less than 1% of the maximum size (i.e. 15 entries). They correspond to isolated or faulty nodes.

3.2 Bidirectional and Unidirectional links

We can distinguish between unidirectional and bidirectional links (RSSI measures are available for one or for both directions). We obtained 16 unidirectional links and 280 bidirectional links.

We define as *link occurrence ratio* the number of appearances of a candidate node in the neighborhood table of a reference node divided by the total number of tables for that reference node. In other words, it represents the percentage of the cases where a link between two nodes was detected and qualified with an RSSI value. We can notice in Figure 2 that a significant number of links (20%) exist less than 1% of the time. By filtering these sets with too small a cardinality, we in particular eliminated all the unidirectional links: their data sets accounted only for 1 to 4 occurrences. Thus, one of our first results is that the testbed did not have any unidirectional links. However, some of the bidirectional links can be asymmetrical (i.e. their quality is different for both directions), as explained in one of the following section.

Unidirectional links may appear when antennas are not perfectly omnidirectional [12], the filters are not well-conceived [13] or when the nodes do not use the same transmission power [14]. Consequently, we can conclude that the Coronis nodes are robust and the hardware is well conceived and industrialized (i.e. different nodes have the same characteristics).

3.3 Filtering data

Since we focus on experimental data, we have to discard ambiguous measures (i.e. possible outliers or impracticable values) to obtain unbiased results. We propose to detect and discard this kind of values.

Formally, we consider that a value is an outlier, if it conforms to the following condition:

$$x < Q1 - 1.5 \cdot IQR \qquad \vee \qquad x > Q3 + 1.5 \cdot IQR \qquad (1)$$

where Q1 represents first quantile of observed data set, Q3 third quantile and IQR difference between them i.e. interquantile range.

Thus, we discard all the values that are single isolated outliers: only one value is extreme, corresponding surely to a transient behavior. On the contrary, multiple consecutive outliers could arise from temporary obstacles (e.g. a delivery truck, a car) for radio propagation, climatological changes (heavy rain that disturbs radio transmissions and increases the humidity measures). Thus, we keep all multiple consecutive outliers. In other words, we consider that the extreme values that last for more than 17 minutes are valid. We will give more attention to multiple consecutive outliers later in the article to infer the main causes and consequences.

After filtering our experimental dataset, we proceed with the analysis.

4. LINK QUALITY

The progress in the radio chip design positively impacted the performance and reliability of WSNs [15]. This motivated us to further investigate the possibility to use the RSSI value as a reliable link quality indicator.

4.1 Radio link Symmetry

We measured the RSSI value in both directions for each radio link (Figure 3). In this graph, we did not remove the links with a very small number of values (as explained in section 3.3) because we aim here at analyzing the reason of their existence.

When the points are close to the diagonal, the links are symmetrical: the quality is identical in both directions. The reader can remark that contrary to the literature, symmetry is predominant.

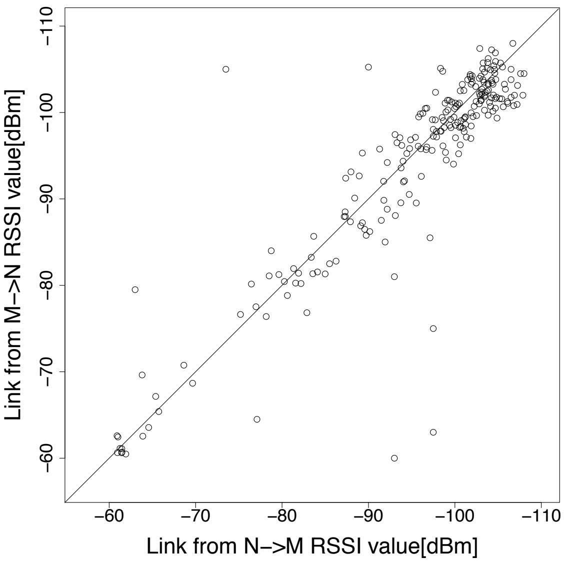

Figure 3: Symmetry of existing bidirectional links

This means that nodes use the same transmission power. Besides, they are also homogenous: the radio hardware behaves identically. For instance, the radio chips of two different radio modules follow the same frequency selectivity (i.e. filters are identical).

Radio links are seldom asymmetrical: these outliers appear for links with a duration less than 1% of total length of the experiment. For these rare cases, the quality in one direction is significantly different, i.e. greater than 10%, sometimes even 55%. This *unbalanced* representation justifies the removing of links with too small cardinalities.

4.2 RSSI distribution

To predict the link behavior with a local and simple metric, the measure should follow a known probability distribution model: we would be able to accurately infer the average quality of the link by analyzing in real-time the measured values.

The Normal (or Gaussian) distribution is extensively used since it models well many natural phenomena, especially for radio propagation (e.g. the Additive White Gaussian Noise). We aim here at verifying if the RSSI measured for each of the existing links follows this distribution.

We applied the Shapiro-Wilk test [16], to the measured RSSI samples. For 92% of the links, the p-value of the Shapiro-Wilk test was significantly less than 0.05 and for the rest barely over this value. This signifies that we need to reject the null-hypothesis, meaning that the RSSI does not follow a normal distribution. This corroborates some indoor results [17] and even in outdoor conditions for LOS radio links, the RSSI does not follow a normal distribution.

We compared also these RSSI samples to other two well-known distributions: Logistic and Cauchy. These distributions are the single ones that permit to have this kind of values (close to a log-normal law but with minor variations).

We used a Kolmogorov-Smirnov test, permitting to compare a well-known distribution to a collection of samples.

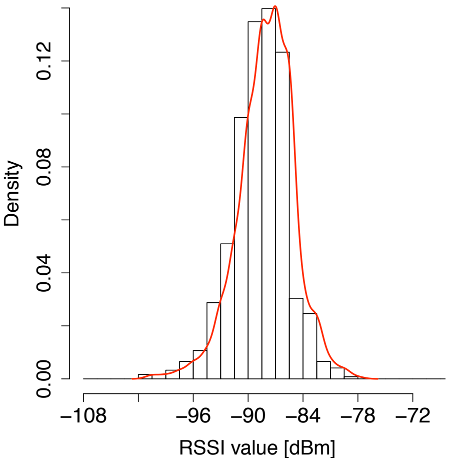

Figure 4: Distribution of RSSI value for one of the representative link

More precisely, a collection of values are generated according to the well-known distribution, with the same cardinality as the set we want to compare to. Then, the Kolmogorov-Smirnov test lets us know if the two collection of values follow the same distribution. For both distribution cases and for all of the links resulting p-value was always close to 0, meaning that we have to drop null-hypothesis i.e. RSSI samples do not follow neither Logistic neither Cauchy distribution. Nevertheless, RSSI distribution that we tried to describe, has bell shape with high central peak but it is slightly skewed to one side (Figure 4). Thus, no well-known distribution permits to have a generic model for such RSSI values.

We now aim at demonstrating that the RSSI of the different links follow the same distribution. Since they do not follow the Normal distribution we have chosen one of the most familiar non-parametric test—the Wilcoxon-Mann-Whitney Test [16]. Now the null-hypothesis is that values from the two independent samples come from the same distribution. The average of p-value for this test was 0.0416 while 89% of values were smaller than 0.05. The null-hypothesis is valid since this p-value is lower than alpha level 0.05. In other words, we conclude that the corresponding pairs of samples do follow the same distribution.

This can be also observed in Figure 5. We plotted the Box Plots for some representative links. Even though median values are not perfectly aligned, we can notice that inter quantile ranges are similar as well as the skewness of data and max/min values.

4.3 RSSI periodicity

We also analyzed the difference in radio link quality during working hours (8am-7pm) and night periods (9pm-6am). Plotting the values we obtained almost the same graph as the one plotted in Figure 3 showing that links did not change their properties during different periods of day. Thus, movements of people and vehicles in the technical park during working hours do not have any significant impact on the RSSI. RSSI is stable and transmissions are quite robust to some environment properties changes.

In other words, the PHY layer in the Wavenis nodes is robust to interferences because it uses frequency hopping. Moreover, the PHY channel is stable.

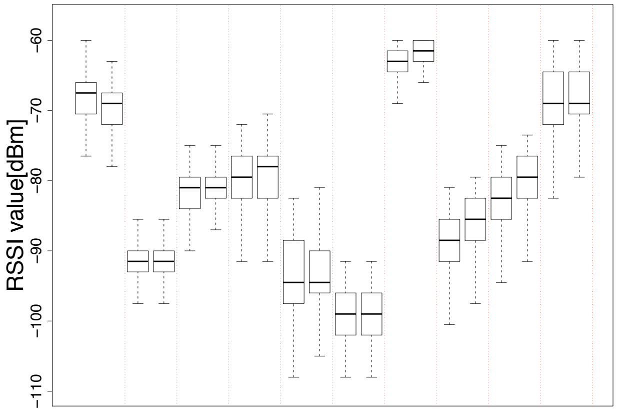

Figure 5: Box plot showing that bidirectional links follow the same distribution (with similar parameter's values

4.4 RSSI vs. Link occurrence ratio

To have a more detailed insight into the link occurrence ratio property shown in Figure 2, we tried to observe whether it can be correlated with the RSSI value.

Figure 6 shows Box Plots for all recognized links in the testbed separated in 4 groups according to the range of their link occurrence ratio without sorting them in ascending order by the same criteria.

Looking at this figure we can notice that there is no evident correlation between RSSI value and link occurrence ratio since Box Plot of RSSI covers whole extent of possible values in different link occurrence ratio ranges. However, we can remark the following points:

- a single RSSI distribution for a particular link does not permit to conclude on the occurrence ratio for this link. Individual conclusions are not possible;
- if we take a closer look at the graph, we can remark that each *category* exhibits different RSSI spreadings. In other words, we could derive a probability of link occurrence ratio for different RSSI values. However, this constitutes a global (and not individual) behavior, i.e. RSSI is not directly a good quality estimator;
- for the first range of link occurrence ratio (1-40%), mean value of the RSSI for all of the links do not pass above -90 dBm. Thus, a poor link means obligatory a low RSSI;
- the largest RSSI values mean in most cases we face stable links.

4.5 RSSI vs. sensor measurements

First, we checked the correlation between the measured humidity and the RSSI. During the experiments, nodes happened to be exposed to humidity levels between 0 and 100% relative humidity (RH). We computed the Pearson's correlation factors [16] for all bidirectional links. In all cases, the value did not exceed 0.5, plus we have neither a negative nor a positive correlation between the two variables. In the

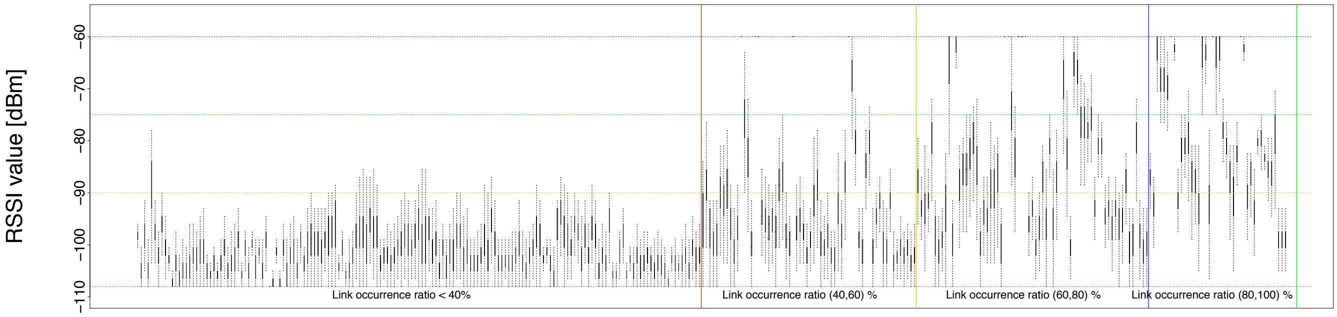

Figure 6: Impact of the RSSI value on link occurrence ratio

same way, the correlation is not significant if we only focus on outdoor radio links.

The second test was an attempt at further exploring the correlation between humidity and RSSI values, but just taking into account the impact of the extreme (maximum) values of humidity measurements. We extracted subset of top 25% of all humidity values (more than 75% of RH) measured during the experiment. Afterward, we computed the difference between the mean value of RSSI for the observed link and the current value of RSSI at the same temporary instant as the measured extreme value of humidity.

If the humidity does not impact the RSSI, we would have an average difference equals to 0. This means that the RSSI value has the same chance to be greater or lower than the average RSSI value. Since we observed this behavior (values are very closed to 0, with a varying sign), we can for sure say that humidity has no impact on the RSSI measurements.

In conclusion, the fast frequency hopping technique is efficient to avoid interferences and frequency-selective fading.

5. NETWORK DYNAMICS

The radio channel is intrinsically unstable, since it is easily influenced by various environmental parameters. This creates a certain dynamism in the network, where links can easily disappear or re-appear. To optimize the performances, the deployed MAC, topology control and routing protocols should self-adapt to changes. We will now focus on the network dynamics to understand how it could further impact the higher layers.

5.1 Neighborhood variation

We first studied the variation in the neighborhood table. The same remarks hold for all the nodes, and we focus here on one node, selected randomly. We plotted in Figure 7 the variation of its neighborhood size.

In the current testbed, there is no `hello` packets because it implements a *all-reactive* solution. When a node wants to transmit a `data packet`, it sends a `RTS`. All its neighbors reply with a `CTS` including the received RSSI. Thus, a node is able to reconstruct the list of its neighbors and the corresponding RSSI.

It varies most of the time with rare stable periods that last at most few samples. We have recognized this behavior as a general trend for all the nodes. This raises the question of whether a proactive approach is the most accurate solution for discovering its neighbors. Indeed, proactive maintenance may result in inefficient routing decisions when choosing non reliable nodes: they can be chosen as next hop because an `hello` was previously received although they will not correctly receive the next data packet. Although RSSI may be stable, the radio link may not be. This result tends to conclude that opportunistic solutions, where the next hop is chosen only when the data packet is transmitted, is more relevant in this environment. Since the next hop is chosen reactively among the nodes that received the data packet, unreliability problem is reduced.

5.2 Link Evolution

Although the neighborhood table changes continuously, a group of stable neighbors may practically exist. In particular, can a large neighbor's stability be correlated with e.g. the RSSI or the distance between the transmitter and the receiver?

We have noticed that stable radio links have one of the following properties:

- a high RSSI value (superior to -75 dBm);
- a pair of nodes within one fifth of the radio range (≈70 m) and having a medium value of RSSI (between -75 dBm and -90 dBm).

By combining distance and RSSI information, we should be able to predict the link stability. Moreover, there was no single case in which neighbors with a high value of RSSI were not among most stable neighbors.

High RSSI could be used as reliable indicator of link stability when using Wavenis chips [10]. Geographical information is a plus to cope with medium RSSI values. Similar observation was made by other authors [15] for different type of radio chips. Nevertheless, there is still substantial free space in order to make tighter conclusions about the link behavior with RSSI in a gray zone (low levels close to the threshold) since it is influenced by various effects (multipath, fading, interference, etc.) for which the impact varies over time and according to the situation.

5.3 Multiple consecutive outlier distribution

As previously stated in section 3.3 we have kept multiple consecutive outliers since they depict transitory effect that influence the quality of radio channel for a short period. An outlier will introduce a bias in the average and median values if we do not discard them.

We aim here at analyzing this phenomenon in a more global way, i.e. for all observed links in the network. Thus, we extracted the Cumulative distribution function (CDF) of

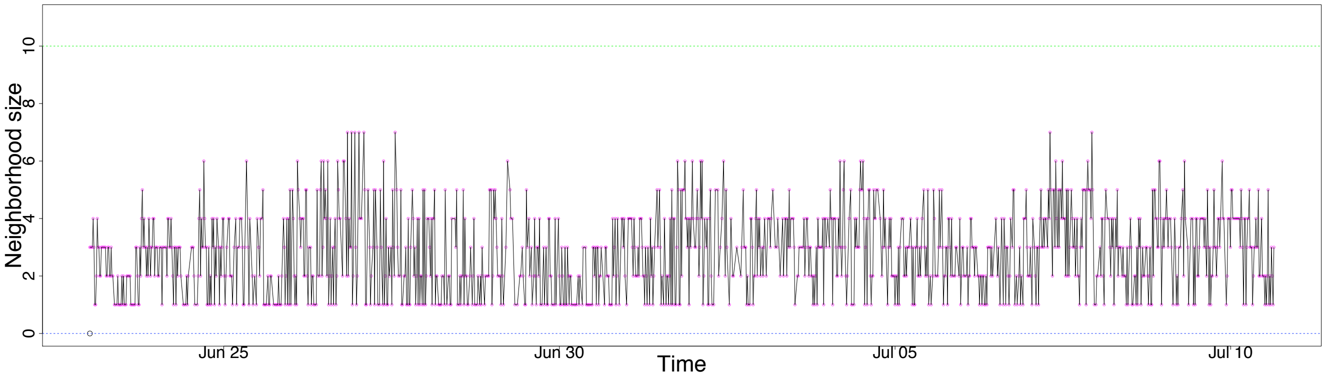

Figure 7: Variation in the neighborhood size for one of the nodes

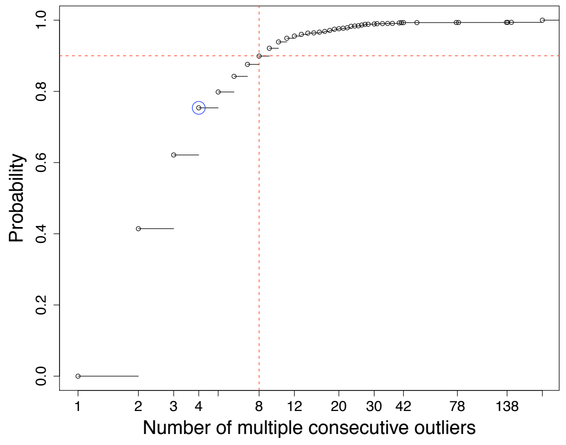

Figure 8: CDF of multiple consecutive outlier

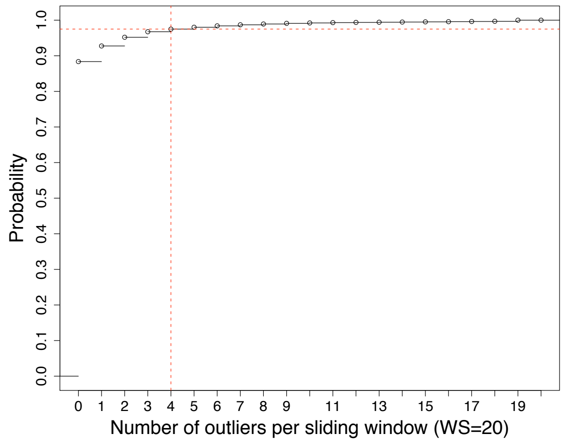

Figure 9: CDF of number of outliers per sliding window of size 20 samples

the number of multiple consecutive outliers. We plotted the results in Figure 8.

In our static testbed we have 4 or less consecutive outliers in 75% of the cases (the blue circle in Figure 8). In the same way, for 90% of the cases, we have 8 or less multiple consecutive outliers (dashed line in the figure). This means that we can consider multiple consecutive outliers lasting up to 8 periods as transitory effects which interrupt stable radio link for short period. Since observing more than 8 consecutive outliers is very rare in a static testbed, we can consider that this phenomenon is related to a permanent topology change in a mobile/changing testbed (e.g. building modification).

Let's still focus on the 4 consecutive outliers case. We have approximately 12% of the samples that last for exactly 4 outliers (75% - 63%)), and only 5% of th samples that last for exactly 5 outliers (80% - 75%). In other words, when a node experiences 4 consecutive outliers, it's more likely that the next sample will be *normal* than it will still be an outlier. We can remark that this observation holds for all cases: we have a strictly larger probability to have k consecutive outliers than k+1. In other words, outliers have a limited impact and the average and mean values would be well-estimated if they are properly detected and discarded (i.e. they will not introduce a large bias).

5.4 Number of outliers in a sliding window

A problem occurs when we want to detect outliers practically. It is almost impossible to keep the whole history in a node memory in order to precisely calculate IQR and to accurately remove outliers. We assume we may only save the last few samples. We chose here to implement a sliding window of 20 samples. Further, we have calculated the Cumulative Distribution Function of the number of outliers per sliding window (Figure 9) to justify our choice.

In 97.5% of cases, we have 4 or less outliers per sliding window (cf. the dashed line). We aim here at limiting the memory consumption while still well estimating the average value to be able to accurately detect the outlier values. Moreover, the method must not be too conservative since testbeds are not static and the environment can change. In particular, the quality should sometimes be re-estimated, even if new values are far from the previous average values.

We propose the following approach to reach this objective. At most 4 slots will be used to store outlier values (yellow fields in Figure 10). These values will be flagged and won't be used to calculate the IQR value (eq. 1) since we consider that these values are *abnormal*. Possibly, a new value could be detected as outlier although 4 values were already flagged outliers. In this case, we un-flag the outlier closest to the

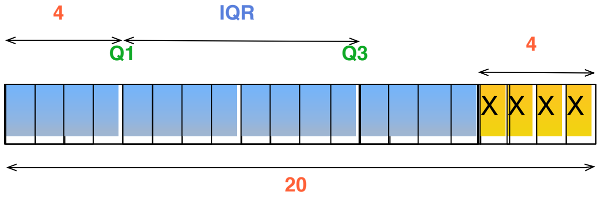

Figure 10: Sliding window example with 4 values flagged as outliers

median value. IQR are updated and possibly the outliers could be considered as *normal* if they lie, after the update, in the correct range.

Let's consider Figure 10. We can see that the extreme 4 values on the right are flagged as outliers and thus are not used to compute IQR, Q1, etc.

Using this approach, we will *smooth* the quality metric and discard inaccurate measures. Moreover, we are also reactive: we efficiently detect changing radio links and update their associated quality metric accordingly.

In conclusion, if each value can be coded in $size_{sample}$, a node only has to reserve $20 * size_{sample}$ bits to compute an accurate average metric.

The reader can note that such a statistical approach could be easily applied to any metric measuring the quality of a radio link.

6. RELATED WORK

After working on theoretic protocol issues and energy savings in WSN, researchers in the last years have taken up experimentally studying the behavior of WSNs [18, 19]. Indeed, simulations are often too simplistic and new problems can arise in experiments that would have not been appeared in simulations. Barren et al. provide guidelines to conceive WSN testbeds [6].

Since experiments require much effort, some platforms aim at simplifying the performance evaluation of protocols and algorithms. Orbitlab [5] deployed a grid of sensors in an isolated environment. Some other approaches have been studied for more finely controlling a testbed and obtaining reproducible experiments [20, 21, 22]. Lei et al. proposed to map real environments into an artificial testbed [23]. All these approaches aim at simplifying the deployment by proposing a controlled environment.

Other approaches aim at deploying an ad hoc testbed for evaluating the performances of one particular protocol, e.g. [5] for AODV and OLSR, [24] for network coding, [25] for a routing metric, etc. Their purpose is specific and researchers have a hard time trying to generalize these results.

Recently, Raman et al. studied the problem of interference in IEEE 802.11 wireless mesh networks and tried to find out if the concept of a radio link is valid in this kind of networks [26]. Their results are experimental, but lead to fundamental and generic results in wireless mesh networks. Although it is focused on mesh networks, it gives an overview of what concerns may raise in Wireless Sensor Networks.

Some authors conducted real-world experiments. In particular, Barren et al. deployed a WSN to obtain meteorological data [27]. These measures will have been later used to derive meteorological models and to predict floods.

Ceriotti et al. deployed a WSN for measuring the movements of the Torre Aquila [28]. The authors provided feedback on the tools an engineer should deploy to design an efficient WSN. Both articles demonstrated that precise measures sometimes cannot be obtained from a single point of measurement. Distributed systems are more robust and permit to obtain rich information. However, they do not provide insights into the WSN environment itself. Thus, we proposed here to fill in this gap.

7. CONCLUSION & PERSPECTIVES

In this paper, we have proposed a way to extend the use and the contribution of implemented testbeds in a urban environment. We have carried out thorough statistical analysis on a collected dataset to obtain an insight on the WSN environment and to emphasize its most distinguished properties.

Our analysis considered the aspects of WSNs such as the link characterization, correlation with environmental parameters (temperature, humidity, and luminosity) as well as network dynamics.

First, we showed that, contrary to the literature, there were no unidirectional links in our observed testbed, and more over that all bidirectional links are highly symmetrical while comparing their mean RSSI values. Furthermore, we have shown that RSSI values don't follow any of the basic distribution models (Normal-Gaussian, Logistic and Cauchy): an accurate distribution still has to be proposed.

Although it is well-known that high humidity may cause decrease in link quality, we have shown that there is no correlation between humidity and RSSI in our experiments. Even extreme maximum values of humidity don't cause significant changes in link quality measurements. This is most probably thanks to the MAC and PHY layers present in the radio chips.

We have highlighted that a proactive approach in neighborhood discovery may cause imprecise routing decisions, favoring the reactive solutions. Besides, although the RSSI exhibits large variations and does not correlate well to the link quality, we were able to characterize the stable links. In particular, high RSSI (more than -75 dBm) or a combination of both distance less than 70 m and RSSI between -75 and -90 dBm permit to conclude that we face stable links.

Finally, we presented dynamic, reactive but still flexible mechanism for detecting and discarding transitory outlier values in measured RSSI value.

Besides, we could now give some advices to researchers who aim at deploying a testbed for characterizing the WSN environments:

- the network should be globally synchronized: we would have been able to compute also average delays, and give upper bounds on delays;
- the Packet Error Rate should be included for both neighborhood discovery packets and data packets in order to obtain a second metric of reliability;

8. ACKNOWLEDGMENTS

The authors would like to thank the personnel of Orange Labs (especially Giyyarpuram Madhusudan) for providing the experimental data.

This work was supported in part by the French Government and the Competitive Clusters Minalogic and System@tic under contracts FUI SensCity and ANR-09-VERS-017.

9. REFERENCES

[1] J.N. Al-Karaki and A.E. Kamal. Routing techniques in wireless sensor networks: a survey. *IEEE Wireless Communication*, 11(6):6 – 28, December 2004.

[2] Giuseppe Anastasi, Marco Conti, Mario Di Francesco, and Andrea Passarella. Energy conservation in wireless sensor networks: A survey. *Ad Hoc Networks*, 7(3):537–568, 2009.

[3] David Cavin, Yoav Sasson, and Andre Schiper. On the accuracy of manet simulators. In *POMC*, pages 38–43, Toulouse, France, October 2002. ACM.

[4] Mineo Takai, Jay Martin, and Rajive Bagrodia. Effects of wireless physical layer modeling in mobile ad hoc networks. In *MobiHoc*, pages 87–94. ACM, Long Beach, CA, USA, 2001.

[5] D. Rastogi, S. Ganu, Y. Zhang, W. Trappe, and C. Graff. A comparative study of aodv and olsr on the orbit tesbed. In *Milcom*, Orlando, Florida, USA, October 2007. IEEE.

[6] Guillermo Barrenetxea, François Ingelrest, Gunnar Schaefer, and Martin Vetterli. The hitchhiker's guide to successful wireless sensor network deployments. In *Sensys*, pages 43–56, Raleigh, USA, November 2008. ACM.

[7] Anukool Lakhina, Konstantina Papagiannaki, Mark Crovella, Christophe Diot, Eric D. Kolaczyk, and Nina Taft. Structural analysis of network traffic flows. In *SIGMETRICS*, pages 61–72, New York, NY, USA, 2004. ACM.

[8] Haakon Ringberg, Augustin Soule, Jennifer Rexford, and Christophe Diot. Sensitivity of pca for traffic anomaly detection. *SIGMETRICS Performance Evaluatation Review*, 35(1):109–120, 2007.

[9] Thomas Watteyne, David Simplot-Ryl, Isabelle Augé-Blum, and Mischa Dohler. On using virtual coordinates for routing in the context of wireless sensor networks. *PIMRC*, 2007.

[10] http://www.wavenis-osa.org.

[11] Thomas Watteyne, Isabelle Augé-Blum, Mischa Dohler, Stéphane Ubéda, and Dominique Barthel. Centroid virtual coordinates - a novel near-shortest path routing paradigm. *Computer Networks.*, 53(10):1697–1711, 2009.

[12] Tereus Scott, Kui Wu, and Daniel Hoffman. Radio propagation patterns in wireless sensor networks: new experimental results. In *IWCMC*, pages 857–862, Vancouver, British Columbia, Canada, 2006. ACM.

[13] R.P. Liu, Z. Rosberg, I.B. Collings, C. Wilson, A.Y. Dong, and S. Jha. Overcoming radio link asymmetry in wireless sensor networks. In *PIMRC*, pages 1–5. IEEE, 2008.

[14] Lifeng Sang, Anish Arora, and Hongwei Zhang. On link asymmetry and one-way estimation in wireless sensor networks. *ACM Transactions on Sensor Networks*, 6(2):1–25, February 2010.

[15] Kannan Srinivasan and Philip Levis. RSSI is under appreciated. In *EmNets*, Cambridge, MA, USA, May 2006.

[16] Peter. J. Huber. *Robust Statistics*. Wiley, 2004.

[17] K. Kaemarungsi. Distribution of wlan received signal strength indication for indoor location determination. In *ISWPC*, Phuket, Thailand, January 2006. IEEE.

[18] Adel Aziz, Alaeddine El Fawal, Jean-Yves Le Boudec, and Patrick Thiran. Aziala-net: Deploying a scalable multi-hop wireless testbed platform for research purposes. In *MobiHoc S3*, New Orleans, USA, May 2009. ACM.

[19] R. Riggio, D. Miorandi, I. Chlamtac, N. Scalabrino, E. Gregori, F. Granelli, and Yuguang Fang. Hardware and software solutions for wireless mesh network testbeds. *IEEE Communications Magazine*, 46(6):156–162, June 2008.

[20] Juan Flynn, Hitesh Tewari, and Donal O'Mahony. A real-time emulation system for ad hoc networks. In *CNDS*, pages 115–120, San Antonio, USA, January 2002.

[21] James T. Kaba and Douglas R. Raichle. Testbed on a desktop: strategies and techniques to support multi-hop manet routing protocol development. In *MOBIHOC*, pages 164–172, Long Beach, USA, October 2001. ACM.

[22] Sagar Sanghani, Timothy X. Brown, Shweta Bhandare, and Sheetalkumar Doshi. EWANT: The emulated wireless ad hoc network testbed. In *WCNC*, New Orleans, USA, March 2003. IEEE.

[23] J. Lei, R. Yates, L. Greenstein, and Hang Liu. Mapping link snrs of real-world wireless networks onto an indoor testbed. *IEEE Transactions on Wireless Communications*, 8(1):157–165, January 2009.

[24] Sachin Katti, Hariharan Rahul, Wenjun Hu, Dina Katabi, Muriel Medard, and Jon Crowcroft. Xors in the air: Practical wireless network coding. *IEEE/ACM Transactions on Networking*, 16(3):497–510, June 2008.

[25] John Bicket, Daniel Aguayo, Sanjit Biswas, and Robert Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *MOBICOM*, Cologne, Germany, August 2005. ACM.

[26] Bhaskaran Raman, Kameswari Chebrolu, Dattatraya Gokhale, and Sayandeep Sen. On the feasibility of the link abstraction in wireless mesh networks. *IEEE Transactions on Networking*, 17(2):528–541, April 2009.

[27] G. Barrenetxea, F. Ingelrest, G. Schaefer, M. Vetterli, O. Couach, and M. Parlange. Sensorscope: Out-of-the-box environmental monitoring. In *IPSN*, St. Louis, Missouri, USA, April 2008. IEEE/ACM.

[28] Matteo Ceriotti, Luca Mottola, Gian Pietro Picco, Amy L. Murphy, Stefan Guna, Michele Corra, Matteo Pozzi, Daniele Zonta, and Paolo Zanon. Monitoring heritage buildings with wireless sensor networks: The torre aquila deployment. In *IPSN*, San Francisco, USA, April 2009. ACM/IEEE.