

Coping with Unreliable Channels: Efficient Link Estimation for Low-Power Wireless Sensor Networks

Andreas Meier, Tobias Rein, Jan Beutel, Lothar Thiele
Computer Engineering and Networks Lab, ETH Zurich
8092 Zurich, Switzerland
a.meier@tik.ee.ethz.ch

Abstract—The dynamic nature of wireless communication and the stringent energy constraints are major challenges for the design of low-power wireless sensor network applications. The link quality of a wireless link is known for its great variability, dependent on the distance between nodes, the antenna’s radiation characteristic, multipath, diffraction, scattering and many more. Especially for indoor and urban deployments, there are numerous factors impacting the wireless channel. In an extensive experimental study contained in the first part of this paper, we show the magnitude of this problem for current Wireless Sensor Networks (WSNs) and that based on the overall connectivity graph of a typical multihop WSN, a large portion of the links actually exhibit very poor characteristics. We present a pattern based estimation technique that allows assessing the quality of a link at startup and as a result to construct an optimal neighbor table right at the beginning using a minimum of resources only. Our estimation technique is superior compared to other approaches where protocols continue to decide on the fly which links to use expending valuable energy both for unnecessary retransmissions and recursive link estimation.

I. INTRODUCTION

Sensor networks have been deployed in many different locations, such as a redwood forest [1], or a potato field [2], having great problems with the reliability of the application’s data transport, often not even yielding a data-delivery rate of fifty percent. These outdoor deployments have to tackle various difficulties, such as the exposure of the nodes; however, when it comes down to the physical level of the transmission, indoor deployments seem [3], [4] to be the hardest to work with due to the high degree of multi-path fading. This results in many unreliable and hardly predictable communication links, which is a crux for indoor deployments, especially for safety critical applications, such as a distributed fire-detection system [5].

The selection of reliably connected neighbors from a wide range of possible neighbors is an essential basis for the application’s sound operation. Commonly the link quality is assessed very brief, if at all, before a neighbor is added to the routing table. The links are then assessed on the fly, i.e. the links’ performance is continuously observed and is replaced (the network is rerouted), if a link does not achieve a required performance. If this is done appropriately, the network will converge to a rather stable route eventually; nevertheless invoking four major difficulties: (1) Especially for low data-rate applications, it takes a long time until a solid

routing table is established. If this process is speed up, i.e. the link estimation is based on less samples, a lot of dynamic is introduced in the network, and likely result in frequent packet losses. (2) Routing on unreliable links results in regular retries on the MAC layer, which does not only waste a lot of energy but also congest the channel unnecessarily. These channel congestions, can have severe effects on neighboring ‘good-quality’ links, letting them degrade due to network internal interference and eventually are being (unnecessarily) replaced also. (3) Switching a communication partner is rather costly, especially if a sophisticated MAC protocol like WiseMAC [6] is being used. These protocols can be tuned to work very energy efficient (concerning the long term); however, finding and setting up a link is very costly due to the required synchronization not even mentioned the network’s overhead for the dissemination of the new route. (4) Energy efficient protocols minimize the overhearing of packets and do not allow keeping track on the surrounding communication traffic (for instance taking track and assess potential neighbor’s packet sequence numbers). So if a neighbor needs to be replaced in the routing table, the costly process of finding and selecting an appropriate neighbor has to start from scratch again.

Selecting only good-quality links right at the beginning of the networks operation would therefore largely increase the application’s performance (energy and time wise), in particular during this start-up. All the current estimators assess the links on the fly and are mostly concerned on how to deal with the limited resource to store the neighbor table. This work, on the other hand, shows how the quality of the link assessment can be improved, in particular how the link quality is to be assessed prior to adding the links to the neighbor table. In a first part we present extensive measurements of the link characteristic in indoor deployments (Section III) and discuss metrics (Section IV) for determining a link’s stability. According to this metrics, we analyze in Section V the links in our networks, showing that the link-quality varies greatly, suggesting for a thorough link assessment. Based on the measured traces, we show in Section IV how the link estimation can be influenced by different parameters, namely the number of packet, the assessment time and the RSSI value, and conclude the paper in Section VII. Furthermore, the data gathered for this study is publicly available online.

II. RELATED WORK

The link quality in wireless sensor networks, i.e. using low power radio devices such as the CC1000, RFM TR1000, has been analyzed in several studies [3], [4], [7]–[10]. Most of these studies focus on the spatial properties of wireless links, based on rather artificial node deployments such as a line or a grid. In contrast, the measurements presented in this work are based on real indoor deployments, in our case for building monitoring. Furthermore, most available studies, were mainly measuring channel quality with byte stream based, low power radios in the sub gigahertz band. Technology however tends towards ample bandwidth, packet based radios, such as the Chipcon CC2420 based on the IEEE 802.15.4 standard, for which only limited studies are available.

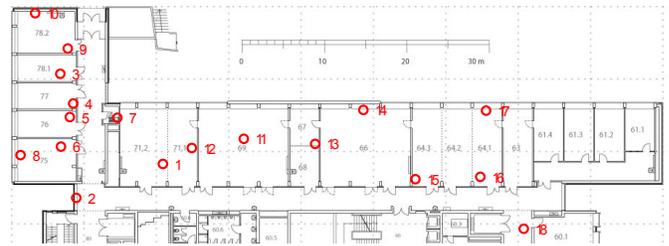
In [3], Zhao et al. (linear topology with a 868 MHz radio device) showed in one of the first empirical studies with low-power radios the existence of a so called *gray area*. That is, there is a fraction of nodes that are only poorly connected, despite there are nodes farther away from the transmitting node that receive almost all packets. Woo et al. [7] (grid, 868 MHz) had similar findings based on a uniform grid over a large, essentially unobstructed indoor space. Reijers et al. [4] (linear, 868 MHz) extensively investigated the gray area and the influence of the environment. They pointed out, that in indoor deployments, due to the multi-path signal delivery, the extent of this gray area is very large, i.e. having more nodes in the grey area than within the proper communication range. Willig et al. analyzed in [8] (linear, 868 MHz) the occurrence of bursts for a sub gigahertz radio, showing that bit errors occur in bursts, packet errors on the other side do not. Cerpa et al. analyzed in [10] (grid, 868 MHz) the temporal properties of wireless links, performed in an indoor office-like setting, based on a grid structure. They studied short term issues, e.g. lagged autocorrelations of individual links, lagged correlations of reverse links, and consecutive same path links. In [9] (random, 2.4 GHz) Srinivasan et al. studied RSSI and LQI of the CC2420 radio. They show that RSSI is usually a better link indicator than LQI due to the latter’s wider range of variance, a result we observed in our testbed as well.

These varying link qualities suggest assessing the link quality of the neighbors that are communicated with. Woo et al. analyze in [7] traditional techniques for the link assessment, pointing out that most of these approaches requires to overhear all neighboring traffic. They propose the so called WMEWMA(t, α) estimator, that does not rely on tracking down the sequence number of the neighboring traffic and mainly elaborate the issue of the resource constraints if many neighbors are available in a possible dynamic environment. However, the link assessment and estimation is not a well addressed topic, despite the various studies on the link quality in WSNs strongly suggesting for elaboration.

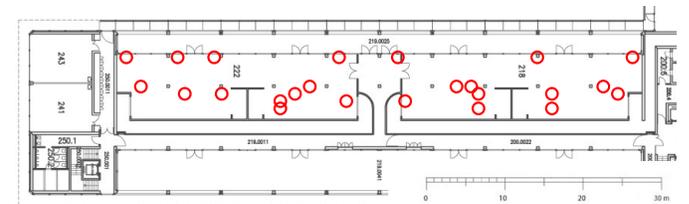
III. EXPERIMENTAL SETUP

As indicated above, most available studies are based on artificial topologies, usually line and grid deployments on a table top, on the ceiling or along a floor, hardly reflecting

the environment an actual WSN would be deployed in. We however wanted to study the effects that will occur in a real deployment, in particular in indoor environments. The link measurements were performed in two different networks, located in different buildings. Network A, as illustrated in Figure 1(a), is deployed in a hostile environment. The nodes are spread over multiple offices, i.e. there are many static obstacles such as walls, but also dynamic ones such as moving people. Furthermore there are many interference possibilities due to 802.11g WLAN access points, a Bluetooth network and devices/machines like elevators, a microwave or a fridge. The Network B, Figure 1(b), on the other hand is deployed in a rather benign, infrastructural part, of a building. Located in the top floor, there is a lot of infrastructure, such as machinery (220/380V), pumps and pipes. In contrast to Network A, there are only few people moving around and the WLAN present is little used, if at all.



(a) Network A having 18 nodes deployed in an office-like scenario. In particular there is a WLAN installed, people are moving, doors are being opened and closed all the time.



(b) Network B has 24 nodes deployed in a less dynamic environment, with only seldomly people entering and moving around.

Figure 1. Network setup.

The measurements were conducted on the Tmote Sky platform, featuring an MSP430 microcontroller and the ZigBee compatible (2.4 GHz) CC2420 radio. For the link measurements one single node was selected transmitting 10’000 packets at 5 Hz frequency (33.3 min), whereas all other nodes in the network were in reception modes, tracing the received packets’ sequence numbers and RSSI/LQI tuples. The transmitting node was then exchanged in a round robin fashion, resulting in a testrun with every node in the network being the transmitting node once, observing all possible links in the network.

In order to see also rare effects in the communication patterns, we performed extensive link measurements, tracing almost 20 Millions packets (Table I). We further used two distinct transmission powers (0 dBm and -5 dBm), which basically doubles the number of analyzed network topologies and links. In addition we conducted the test during workdays

Table I
NETWORK CHARACTERISTICS.

Network	A	B
Number of Nodes	18	24
Number of Links	172	390
Received Packets	7'936'690	7'171'839
Traced Packets	10'430'000	8'280'000
Average PRR	76.1%	86.6%

and nights in order to see the impact having people moving and the increased communication of the WLAN hotspots. The measurements were all conducted on one single channel (26), in order to not further increase the number of uncertainties.

IV. LINK BEHAVIOR AND METRICS

Wireless links are traditionally qualified based on the average packet reception rate (PRR) and is being used in almost all the studies discussed in Section II. This metric reflects the link's overall performance over the measured period; for instance, if 100 packets are sent, but only 95 of them received, a PRR of 0.95 results. The PRR, giving a first impression on the link's performance, lacks the information of the temporal characteristic, i.e. whether the link quality is stable over time.

Figure 2 shows sample traces of such a link's temporal behavior, we observed in Network A: The leftmost plot represents a link that is rather stable over time, i.e. the PRR varies slightly only. With the other three plots this is different. The second plot shows a distinct dropout lasting for 3.5 minutes, the third showing a degrading behavior and the rightmost a suddenly increased link performance after showing a bad link quality for 15 minutes. A rather surprising observation can be made if the RSSI value is taken into account. The stable leftmost link shows the expected very stable RSSI value. On the other hand this holds as well for the rightmost instable link; both even having a very similar average RSSI value. Hence neither the PRR nor the RSSI link metric can reflect the instabilities. This however is of utmost importance for the network protocol, which preferably routes on stable links in order to minimize expensive retransmissions and changes of the topology.

In order to quantify a link's temporal performance, Cerpa et al. [10] proposed to use the so called *required number of packet* (RNP) metric, representing the number of required (re)transmissions upon a successful reception. While the RNP metric being very useful to predict the retransmission overhead for a distinct link of the routing protocol, it showed to lack in providing a good representation of the link stability. Instead we propose to use the standard deviation σ_m of the *temporal* PRR, i.e. aggregating m packets using a sliding window. The degree of the link's stability is reflected in the comparison of the link's standard deviation with the one of the (stochastic) Bernoulli Process σ_m^{rand} and is further detailed and analyzed in Section V-A. Furthermore, we show that the σ_m metric can conveniently be used to distinguish between stable and unstable links.

All the metrics, namely the PRR, RNP and our σ_m , represent a link with a single number only, enabling a direct

comparison of links. However, rare corner cases of a link's behavior will evidently average out. In particular interest for a network protocol is the information about the possibility of a total temporal link failure. Especially the retransmission of packets can be designed more efficient if the behavior such bursts of packet losses are known. Willig et al. [8] analyzed these burst and proposed describing them in terms of the relative frequency of a certain burst size. That is, compared to all bursts, how many consist of a single packet, how many consist of two and so on. The occurrence of packet bursts is analyzed and discussed in Section V-B.

V. STATISTICAL LINK ANALYSIS

The extensive measurements in our network showed, as indicated in Figure 2, that a network and MAC protocol has to deal with unpredictable and unreliable links. Our results have shown that the occurrence of such unpredictable links depends a lot on the environment the network is deployed in. In particular we observed more than 10% of unstable links in Network A, whereas Network B only showed 1.9%. Finding the reason for these sudden changes in the link quality showed to be impossible. For instance, the 2nd and the rightmost plot in Figure 2 were measured in the middle of the night where no people were moving or having doors being opened or closed. Also the WLAN can hardly be blamed, as there is most likely largely reduced WLAN communication during the night. All in all there is no solid explanation possible for the observed link instabilities but they are analyzed in detail subsequently.

Furthermore, the distributed nature of a multi-hop WSN results in having many neighbors that are only poorly connected, i.e. having a PRR far less than 80%. Such unpredictable and poor links make the task for a communication protocol rather challenging and demand for a thorough assessment of the link quality: this not only during the operation but also before setting up the routes, allowing to minimize the resource-ravaging dynamics of the network topology in the early stage of the operation.

A. Link Stability – σ_m Metric

The trace of a link measurement can conveniently be represented by a series of '1's (packet received) and '0's (packet not received). For a perfectly stable link, this sequence is a Bernoulli process, i.e. with independent random variables having each the probability P_r (equal the PRR of a stable link) of being one. In order to map the stochastic process a temporal PRR, m packets are aggregated leading to a binomial distribution $P_m(P_r) = \frac{1}{m} \text{Bin}(P_r)$. The standard deviation σ_m^{rand} of a stochastic and perfectly stable link can then be expressed as

$$\sigma_m^{\text{rand}}(P_r) = \sqrt{\text{Var}(P_m(P_r))} = \sqrt{\frac{1}{m} P_r(1 - P_r)}. \quad (1)$$

For the measured traces on the other hand, the standard deviation can directly be derived from the variance of the temporal PRR, also aggregating m packets:

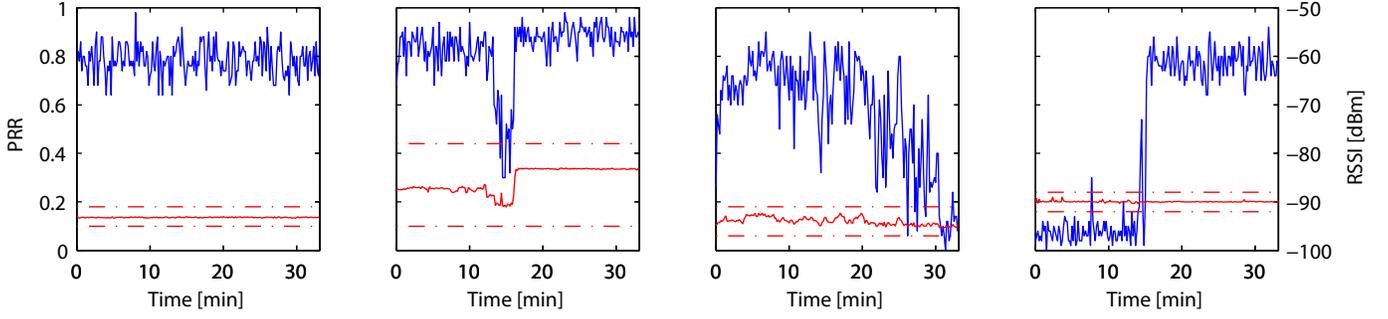


Figure 2. Trace of the PRR (upper curve) and the according (mean, min, max) RSSI (lower curve) of four sample links: Most links in the network show a rather stable link quality over time as shown in the leftmost plot (89.6% in network A and 98.1% in network B). There are also links showing unstable behavior, such as dropouts (2nd), degradation (3rd) or even step behavior (4th). Surprisingly only the 2nd plot shows a distinct change of the RSSI value.

$$\sigma_m^{\text{trace}} = \sqrt{\text{Var}(\text{trace})} \quad (2)$$

Figure 3 shows the standard deviation of all the traces measured in network A and B. The standard deviation of the Bernoulli process (stable link) is plotted by the straight line, whereas all the different traces taken in the two networks are represented by a circle. The difference between the two networks is very distinct, showing a lot of links with a largely increased standard deviation, i.e. increased instability, in network A but only a few in network B. This observation can also be made by manually inspecting the traces of the two networks.

A stable link should have a constant PRR and therefore having a similar standard deviation than the Bernoulli process with the same average PRR. An unstable link on the other hand is expected to show an increased standard deviation, and hence the *stability factor* $\gamma_m = \sigma_m^{\text{trace}} / \sigma_m^{\text{rand}}$ can be used to determine the degree of the link's stability:

$$\gamma_m = \sqrt{\frac{\text{Var}(\text{trace})}{m} P_r (1 - P_r)}, \text{ where } P_r = \text{PRR}(\text{trace}) \quad (3)$$

The traces shown in Figure 2 show an increased instability going from the left to the right: whereas the leftmost trace is rather stable, the step behavior of the rightmost links is definitely not. This behavior is reflected by the stability factor γ_{50} showing an increasing value for the four traces, namely 1.11, 2.33, 3.23 and 4.84, going from the left to right.

The index m of the stability factor γ_m represents the granularity of the link instabilities: Whereas a small value m shows very short time fluctuations of the link, a larger value of m allows detecting long term instabilities. In order to detect whether a link is stable, it should neither show short nor long time fluctuations.

Based on a cumbersome manual inspection of all traces of our extensive link measurements, we were able to define a set of constraints practical for the link stability according to Definition 1. In particular $\gamma_{100} > 3.0$ showed to detect links with a lot of short time fluctuations, and is represented by the dashed line in Figure 3. Long term instabilities on the other

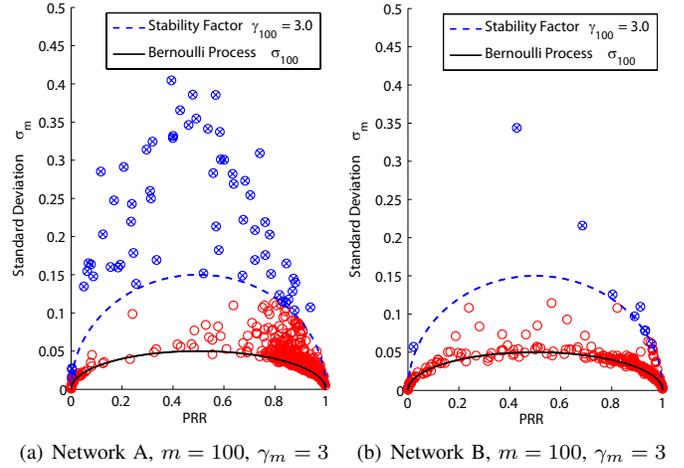


Figure 3. Link Stability: The solid line represents the standard deviation of a stable link; the dashed line indicates the boundary of stable links. Each circle represents an instance of a measurement traces. Network A shows significantly more unstable links.

hand, can be detected using a second criteria $\gamma_{500} > 4.8$. Combining the two criteria finally allows distinguishing the stable from the unstable links.

Definition 1. Stable Link Criteria:

$$\text{Stable Link} = \gamma_{m_i} < \alpha_i \quad \forall i \in [1, I]$$

Figure 3 indicated that network A, deployed in a rather malicious environment, shows much more instabilities than network B, deployed in a benign environment. Using the two constraints for short and long-term instabilities allows quantifying this observation and is presented in Table II. Most notable: a routing protocol in network A has to deal with 10.4% unstable links whereas network B shows only 1.9% such links. This inherently suggests that the parameter optimization of the topology maintenance depends largely on the network's deployment site.

B. Burst Size

Analyzing the occurrence of bursts was proposed by Willig [8]. Their measurements are based on a setup with

Table II
LINK STABILITY: NETWORK A SHOWS AN INCREASED NUMBER OF UNSTABLE LINKS.

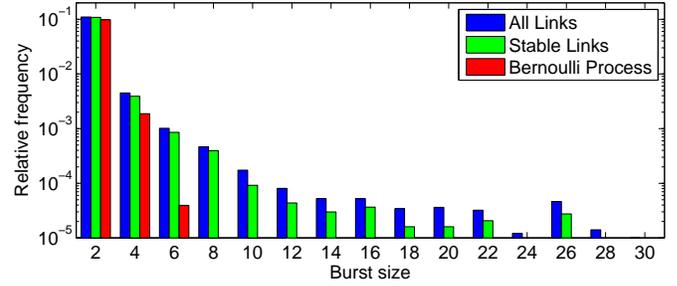
Network	A	B
Number of traces	656	828
Number of packets per trace	10'000	10'000
Packet frequency	5 Hz	5 Hz
Short time instabilities ($\gamma_{100} > 3$)	9.6%	1.1%
Long term instabilities ($\gamma_{500} > 4.8$)	9.4%	1.9%
Fraction of stable links	89.6%	98.1%
Fraction of unstable links	10.4%	1.9%

10 receivers (RFM TR1001), lined up on plank having a node distance of 30 cm (3 m in total). They measured the burst occurrence with the same packet frequency of 5 Hz than we did in our tests allowing a comparison of the results. However, their testbed was rather small and they had no nodes close at the reception range, which resulted in a minimal PRR of 0.908. In order to compare their results we selected our links according to this minimal PRR of their testbed. Not surprising, our testbeds, located in a much more hostile environment, show more frequent bursts, having 'only' 91.2%/95.0% (Network A/B) 'single' bursts compared to the 98.4% in [8]. In particular Network B, not having distinct interfering networks close by, also shows an increased number of bursts. This suggests that the 868 MHz frequency, used by the RFM TR1001, is much less susceptible to burst errors than the CC2420 2.4 GHz frequency.

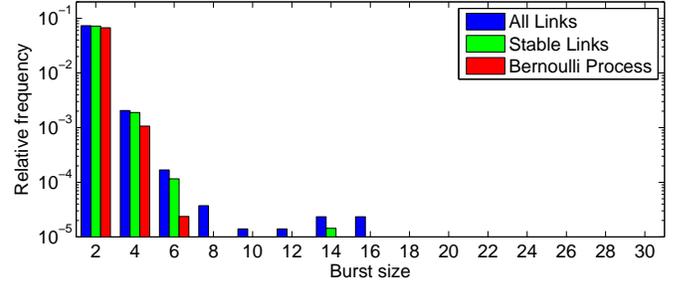
Packet bursts become even more frequent in our testbeds, especially in Network A, if we consider all links with a PRR > 0.8 , as illustrated in Figure 4. Even though more than 80% of all packet bursts consist of a single packet failure, our hostile Network A shows packet bursts with up to 30 subsequent packet misses, i.e. six seconds with the impossibility for communication. Network B on the other hand shows a rather similar behavior than the Bernoulli process, i.e. a link with totally independent packet transmission failures, especially if only the stable links are accounted. This correlation of the packet-failure bursts and the link stability does not hold for Network A. This indicates that the dynamics in Network A, in particular the moving people and the interference of the WLAN, is likely to introduce frequent short-time outages. Hence an aggressive retransmission strategy is likely to waste a lot of energy and might even result in unnecessary topology changes due to the subsequent unsuccessful retries.

C. Long Term Analysis

We repeated the link measurements of the rather unstable Network A after a week and a month respectively, allowing to analyze the long term behavior of the links. About 68% of the links have a standard deviation of the PRR of these three measurements of less than 0.01; 76% of all links less than 0.04. However, almost one fourth of all links have a variance of 0.04 or greater. In other words, these links vary more than 20% referring to the PRR. The causes for this behavior have not been determined yet. The instability discussed in the previous paragraph might have an effect on this result, but there exist



(a) Network A shows a distinct amount of packet-failure bursts with a size of up to 30 packets, i.e. 5 seconds without the possibility for a successful packet transmission.



(b) Compared to Network A, Network B exhibits much fewer concurrent packet failures; not even on the unstable links.

Figure 4. Burst size of packet failures for links with PRR > 0.8

also links having stable behavior in all the three tests, but with a distinct difference in the PRR.

VI. EFFICIENT LINK ASSESSMENT & ESTIMATION

It has been shown previously that a node is likely to have neighbors that are only poorly connected; i.e. only a few of the packets sent are received. For a low-power routing protocol however it is essential that only good quality links are selected, allowing for a reliable and efficient operation. Typical routing protocols usually select a 'random' neighbor for their routing and assess the link on the fly, i.e. during operation, exchanging nodes in the routing table with other (random) neighbors when conditions change. This approach leads to a lot of inefficient adaptation of the network topology, possibly also oscillations, in particular during the start-up of the network.

In contrast to this 'best practice' we suggest to assess a nodes link quality prior to addition to the routing table, only adding it if reasonably well connected. Of course, the continuous link assessment should still take place, possibly profiting from the novel link assessment scheme executed at startup. The question answered in this section is how such a scheme can be performed most efficiently identifying the best links with a high probability. In particular we are analyzing the impact of the packet number and the time for the link assessment and further show that the use of the RSSI value allows for a further increased performance.

A. Evaluation Method

The evaluation of the different strategies for the link assessment are based on 1'448 measured traces as presented in

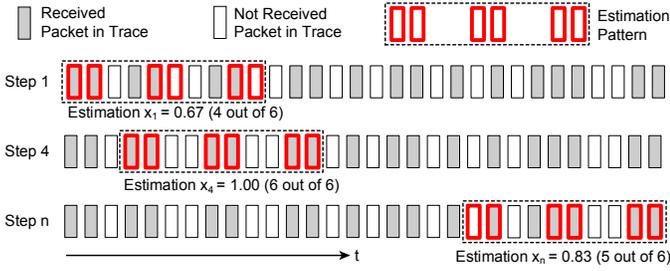


Figure 5. Based on the measured traces a sliding window is used to analyze the performance of the estimation pattern.

Table II, inherently providing information on a trace's average packet reception rate PRR X . The link assessment estimates the PRR based on a so called *estimation strategy*, being either a static *pattern* or an adaptive one, using RSSI information, as illustrated in Figure 5. In this example a static estimation pattern selects the packet numbers 1, 2, 5, 6, 9 and 10 from the trace and estimates the link quality based on these six packets' success rate. By using a sliding window technique, this single pattern will provide $n \simeq 10^4$ link estimations ($x_1 \dots x_n$) for every evaluated trace. For instance, starting the pattern on the first packet results in $x_1 = 4/6$ received packets, whereas starting from the fourth packet results in receiving all packets $x_4 = 6/6$. The n estimations $x_1 \dots x_n$ are compared with the link's 'real' performance X by calculating the variance, where a smaller variance denotes a more precise estimation:

$$\sigma^2 = \frac{1}{n} \sum_{i=1}^n (X - x_i)^2 \quad (4)$$

Evaluating this estimation pattern for all $N = 1448$ traces will result in a variance vector V with N elements.

$$V = (\sigma_{\text{Link1}}^2, \sigma_{\text{Link2}}^2, \dots, \sigma_{\text{LinkN}}^2)' \quad (5)$$

For the comparison of different estimation strategies \mathcal{A} and \mathcal{B} , the two corresponding variance vectors ($V^{\mathcal{A}}$, $V^{\mathcal{B}}$) requires standardization of the corresponding elements. This is necessary to weight all evaluated traces of the link estimation the same; otherwise the evaluation would mainly be based on the links with the higher variance.

$$v_i^{\mathcal{A}} = \sigma_i^{\mathcal{A}} / v_i^{\mathcal{A}} = 1, \quad v_i^{\mathcal{B}} = \sigma_i^{\mathcal{B}} / v_i^{\mathcal{A}} \quad \forall i \in [1, N] \quad (6)$$

The two estimation strategies can then be compared by setting the average of both vectors ($V' = \bar{V}^{\mathcal{B}} / \bar{V}^{\mathcal{A}} = \bar{V}^{\mathcal{B}}$) into proportion, e.g. $V' < 1$ denoting estimation strategy \mathcal{B} being more precise and therefore superior.

Despite this offline evaluation of the link assessment, the results would have been exactly the same assessing the links online. However, this offline approach allows for an elaborate analysis and allows for a fair comparison of different link assessment strategies, all using the very same traces.

B. Static Estimation Pattern

We start with a rather simple estimation pattern, which is going to be the reference for the more elaborate estimation strategies, i.e. is used for the standardization (6) ($\bar{V}^{\text{SA}} = 1$):

Standard Pattern (SA): The standard pattern consecutively sends 20 packets a frequency of 5 Hz. The estimated link quality is the PRR over these 20 packets.

A solid link estimation should not only separate the wheat from the chaff, i.e. differentiating between high and low-quality links, but also give an accurate estimation of the link quality, e.g. whether a PRR of 85% or 95% can be expected.

The standard pattern (SA) reliably performs the task of separating the good from the bad links: The links with a PRR < 0.5 are estimated with a negligible probability of 0.02% of having a PRR > 0.8 , never estimating a PRR > 0.9 . High-quality links with a PRR > 0.9 , are estimated with a probability for 1.84% of having a PRR < 0.8 and only 0.02% of having a PRR < 0.5 .

The performance of estimating high-quality links with a finer granularity is shown in Figure 6(a), depicting the detailed assessment for the links with a PRR > 0.8 in steps of 0.05. The links with a very high PRR > 0.95 are assessed with a probability of 71.1%. For the links with a PRR lower than 0.95, the chances for an exact link estimation drops significantly to less than one third. On the other hand, chances for a distinct misjudgment, i.e. more than 10% difference between the evaluated and estimated PRR, is also less than one third.

The SA solidly distinguishes between high and low-quality links. However, the task of assessing the links with a fine granularity is much harder to tackle and possibilities for improving the performance are discussed subsequently.

Pattern T-X: These patterns estimate the link with the PRR out of 20 packets. The value X indicates the packet interval and therefore determines the estimation period. For instance T-10s traces a packet every 10 sec, requiring 190 sec for the estimation. T-0.2s is equal to the SA.

The evaluation of the pattern T-X is shown in Figure 7(a). It can be seen that a lower packet frequency improves the estimation; however, the improvement gets insignificant for packet intervals longer than 5 sec. A relationship to the burst size of the packet losses evaluated in Section V-B is very likely. If such a burst occurs during the estimation, and the estimation is mainly based on packets of this burst and the resulting estimations will likely be incorrect. Therefore, a packet interval that is longer than the majority of the burst sizes of only a few seconds improves the quality of the estimation. The pattern T-10s is shown in Figure 6(b) exhibiting a slightly increased performance compared to the SA.

We also evaluated other, more complicated, patterns. For instance we took 5 packets with a frequency of 1 packet/sec, then we waited for 20 seconds and we took another 5 packets with a frequency of 1 packet/sec and so one until we acquired

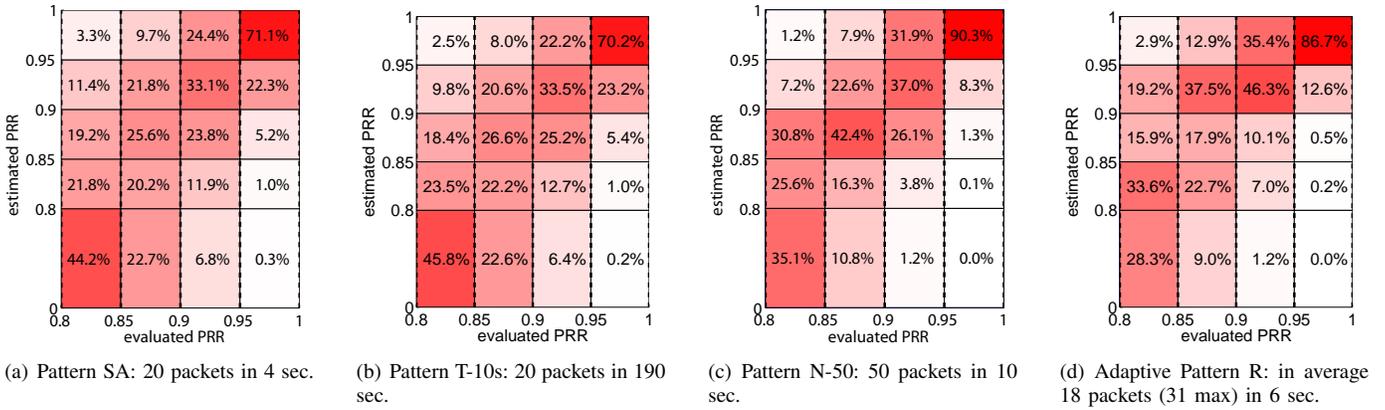


Figure 6. Detailed comparison of different estimation strategies for $PRR > 0.8$

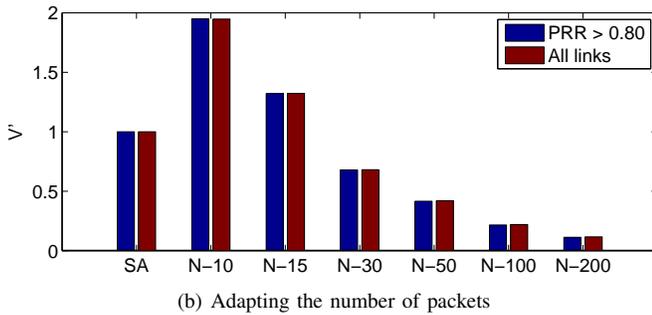
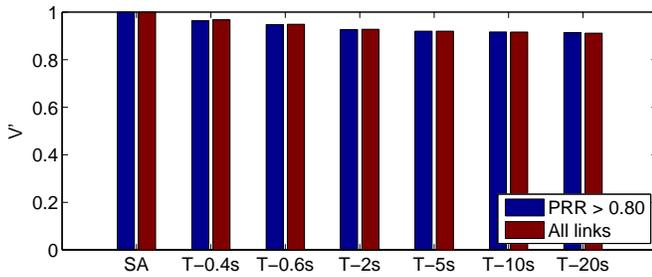


Figure 7. Comparison of different estimation-strategy parameters.

20 packets. These patterns improved the estimation in the same scale as the algorithms T-X. Hence, the basic characteristic for the pattern's efficiency is the temporal extension of the link estimation.

Pattern N-X: These patterns consecutively send X packets with a frequency of 5 Hz. The estimation is performed by calculating the PRR over these X packets. The algorithm N-20 is equal to the SA.

The number of packets has a clear impact on the quality of the estimation as detailed in Figure 7(b). The relation between the number of packets X and the resulting value \bar{V} (compared with the SA) is almost indirectly proportional. This means that if we double the number of packets, the variance of the estimation is divided by two. This result however needs to be quantified with respect to the power and time requirements:

Table III

STATISTICS ON THE PRR FOR DIFFERENT RSSI THRESHOLDS: EVEN LINKS WITH A LOW RSSI CAN HAVE A VERY HIGH PRR.

RSSI Range	Avg PRR	$\sigma(\text{Avg PRR})$	Max PRR
-70...0 dBm	98.2%	0.022	99.9%
-75...-70 dBm	94.6%	0.045	99.0%
-80...-75 dBm	91.5%	0.088	98.8%
-85...-80 dBm	88.7%	0.076	98.6%

For instance, the pattern N-50 requires 2.5 times more energy and time than the SA, but shows a distinct improvement for assessing the high quality links, as depicted in Figure 6(c). In particular, 90.3% of the links with a $PRR > 0.95$ are estimated correctly, in contrast to the 71.1% of the SA.

C. Adaptive Estimation

As indicated in Table III, and previously proposed in [9], the RSSI value allows refining the link estimation. The table shows, that for high RSSI values, i.e. $RSSI > -70\text{dBm}$, the average PRR is more than 98%, showing only a small standard deviation, and hence allows to detect very good links. On the other hand, the RSSI indicator cannot be used to determine whether a link is bad, since there are also links with a very high PRR having a low RSSI value only. The following adaptive pattern R shows how this additional information can be used for an increased performance of the link assessment.

Adaptive Pattern R: The algorithm R is based on a pattern that sends 1 packet/sec up to maximum of 31 packets. The algorithm is divided into three rounds: (1) If the first packet is received having an $RSSI > -70\text{dBm}$ a $PRR = 0.98$ is estimated and the algorithm stops. (2) Another 15 packets are sent. The algorithm stops returning the PRR of the first 16 packets if this $PRR > 0.9$ AND $\min(RSSI) > -80\text{dBm}$. (3) Another 15 packets are sent and the PRR is returned.

In our evaluation this adaptive algorithm with our data, sent an average number of 18 packets. In particular 36% of the estimations stopped after the first round, 13% after the second, whereas 51% of the links required sending all 31 packets

for the estimation. With less packets sent in average, this adaptive algorithm achieves a significantly better assessment ($V' = 0.75$) than the SA pattern and is further detailed in Figure 6(d). The overall performance is slightly below the one of the N-30 pattern ($V' = 0.68$). This indicates that for many good links a lot of energy and time can be saved by quickly assessing them, still pursuing a thorough assessment for the less unpredictable links.

We also analyzed the impact of the LQI for an increased performance for the link estimation and extended the algorithm R in various ways. However, none of these extensions could achieve reasonably better results. Furthermore, tweaking the parameters of algorithm R for an optimized performance is not recommended, since this optimization depends mostly on the environment the target network is being deployed in.

A reasonable value for a minimal link quality (MLQ) for a neighbor in the routing table is a $PRR \geq 0.9$. Randomly selecting a neighbor for setting up a route results in 44.6% of the cases of a later change in the topology since the link does not provide this MLQ. On the other hand, if the adaptive algorithm R is being used for a prior link assessment, this percentage is reduced to 14.8%, i.e. resulting in 3.0 times less network topology changes during the operation. This coming for an average cost of sending 18 packets and only having a 0.4% possibility of wrongly estimating such a MLQ neighbor with a bad link quality of $PRR < 0.8$.

VII. CONCLUSION

With this work we have gained significant insight into the channel characteristics and link behavior of a 2.4 GHz indoor environment typical for wireless sensor networks. The datasets produced by long-term all-to-all channel quality measurements are a substantial reference for future work and provide a sound base for statistical analysis. In particular we presented a novel metric σ_m that allows determining the degree of a link's stability and allows for distinguishing between stable and instable links.

Based on the observations made we have developed a link estimation scheme that allows to efficiently discriminate bad from good links and further specify a detailed link quality metric that can be computed for all links of and above a given threshold quality. This estimation scheme is especially beneficial when used at the startup of a system where the network topology yet needs to be defined simplifying the startup of common topology control and routing protocols. However other use cases are possible as well. In contrast to other approaches the scheme presented is highly effective and uses only limited resources. The patterns developed for link estimation allow selecting the right pattern for different performance goals, such as time and energy constraints. Specifically, the N-50 pattern uses 2.5x more time (and energy) than the SA pattern (when sending 20 packets) but achieves a performance increase of 20% (71.1% vs. 90.3%) which in practice reduces false positive estimations by two thirds. If the duration and resource expenses for the N-50 pattern are too high for a certain application an adaptive estimation strategy

using repetitive packets for link estimation but terminating on a given goal of the packet reception rate (PRR) can be used at the cost of a slight performance decrease.

First results of performing a link assessment prior to addition of the links to the routing table have shown to significantly minimize the number of neighbor changes required subsequently. In particular the dynamics in the network topology during the application's start-up is minimized allowing for a much faster start-up while largely reducing the costly routing overhead.

ACKNOWLEDGMENT

The work presented in this paper was supported by CTI grant number 8222.1 and the National Competence Center in Research on Mobile Information and Communication Systems (NCCR-MICS), a center supported by the Swiss National Science Foundation under grant number 5005-67322. We thank Michael Amrein for the fruitful discussions about statistical analysis and to Matthias Woehrle for proofreading draft versions of this paper.

REFERENCES

- [1] G. Tolle, J. Polastre, R. Szewczyk, D. Culler, N. Turner, K. Tu, S. Burgess, T. Dawson, P. Buonadonna, D. Gay, and W. Hong, "A macroscope in the redwoods," in *Proc. 3rd ACM Conf. Embedded Networked Sensor Systems (SenSys 2005)*, (New York, NY, USA), pp. 51–63, ACM Press, 2005.
- [2] K. Langendoen, A. Baggio, and O. Visser, "Murphy loves potatoes: Experiences from a pilot sensor network deployment in precision agriculture," in *Proc. 20th Int'l Parallel and Distributed Processing Symposium (IPDPS 2006)*, pp. 8–15, IEEE, Piscataway, NJ, Apr. 2006. 14th Int'l Workshop Parallel and Distributed Real-Time Systems (WPDRTS 2006).
- [3] J. Zhao and R. Govindan, "Understanding packet delivery performance in dense wireless sensor networks," in *First Int'l Workshop on Embedded Software (EMSOFT 2001)*, pp. 1–13, 2003.
- [4] N. Reijers, G. Halkes, and K. Langendoen, "Link Layer Measurements in Sensor Networks," in *Proc. 1st Int'l Conf. on Mobile Ad-hoc and Sensor Systems (MASS '04)*, Oct. 2004.
- [5] M. Strasser, A. Meier, K. Langendoen, and P. Blum, "Dwarf: Delay-aware robust forwarding for energy-constrained wireless sensor networks," in *Proceedings of the 3rd IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS 2007)*, pp. 64–81, June 2007.
- [6] A. El-Hoiydi and J. Decotignie, "WiseMAC: An ultra low power MAC protocol for multi-hop wireless sensor networks," in *Proc. 1st Int'l Workshop Algorithmic Aspects of Wireless Sensor Networks (ALGOSENSORS 2004)* (S. Nikolettseas and J. Rolim, eds.), vol. 3121 of *Lecture Notes in Computer Science*, pp. 18–31, Springer, Berlin, June 2004.
- [7] A. Woo, T. Tong, and D. Culler, "Taming the underlying challenges of reliable multihop routing in sensor networks," in *Proc. 1st ACM Conf. Embedded Networked Sensor Systems (SenSys 2003)*, (New York, NY, USA), pp. 14–27, ACM Press, 2003.
- [8] A. Willig and R. Mutschke, "Results of bit error measurements with sensor nodes and casuistic consequences for design of energy-efficient error control schemes," in *Proc. 3rd European Workshop on Sensor Networks (EWSN 2006)*, vol. 3868 of *Lecture Notes in Computer Science*, pp. 310–325, Springer, Berlin, Feb. 2006.
- [9] K. Srinivasan and P. Levis, "RSSI is under appreciated," in *Proc. 3rd IEEE Workshop on Embedded Networked Sensors (EmNetS-III)*, May 2006.
- [10] A. Cerpa, J. L. Wong, M. Potkonjak, and D. Estrin, "Temporal properties of low power wireless links: modeling and implications on multi-hop routing," in *Proc. 6th ACM Int'l Symp. Mobile Ad Hoc Networking and Computing (MobiHoc 2005)*, (New York, NY, USA), pp. 414–425, ACM Press, 2005.