

The Triangle Metric: Fast Link Quality Estimation for Mobile Wireless Sensor Networks

(Invited Paper)

Carlo Alberto Boano*, Marco Antonio Zúñiga†, Thiemo Voigt‡, Andreas Willig§, and Kay Römer*

*Institute of Computer Engineering, University of Lübeck, Germany
{cboano, roemer}@iti.uni-luebeck.de

†Digital Enterprise Research Institute, Galway, Ireland
marco.zuniga@deri.org

‡Swedish Institute of Computer Science, Kista, Sweden
thiemo@sics.se

§University of Canterbury, Christchurch, New Zealand
andreas.willig@canterbury.ac.nz

Abstract—Link quality estimation is a thorny problem in wireless sensor networks, because its accuracy affects the design and the efficiency of networking protocols and applications.

Especially in the context of low-power wireless, estimating the link quality poses a sort of *catch-22* dilemma, whereby a large number of packet samples are required to accurately estimate a channel, but only a few samples should be used due to limited energy resources. This paradox becomes even more severe in mobile wireless sensor networks, since the high variability of the medium imposes even stricter constraints on the timing in which the estimation has to be carried out.

In this paper we propose the *Triangle Metric*, a metric that combines geometrically the information of PRR, LQI, and SNR into a robust estimator that guarantees a fast and reliable assessment of the link quality. Our evaluation shows that the triangle metric can identify the quality of links using as few as 10 packet samples, making it an eligible solution for highly mobile sensor networks.

I. INTRODUCTION

Wireless communication links are well-known to be error-prone and time-varying, major reasons being path loss, shadowing, fading, or external interference. Observing the behaviour of a wireless link and using the results to derive predictions about the future link quality is fundamental to adjust protocols and radio parameters to satisfy the communication requirements in an energy-efficient fashion.

In general terms, the link quality estimation problem can be succinctly described as follows: during a time window $[t_0, t_1]$, called *observation window*, a node collects information from received packets in order to predict the delivery capacity of the link over a certain time horizon $[t_1, t_2]$, called *prediction window*. In wireless sensor networks there is a tradeoff: larger observation windows improve the accuracy of the prediction, but also increase the consumption of the highly constrained energy budget. A link estimator should be accurate, agile, efficient, and should minimize both memory requirements and traffic overhead.

In this work, our goal is to *classify* links into distinct categories according to their quality. We focus on estimating the link quality in mobile wireless sensor networks using IEEE 802.15.4-compliant transceivers [1] such as the Chipcon cc2420 [2]. This setting imposes rigid constraints on the size of the observation window, and on the type of samples available as input to the prediction scheme.

The size of the observation window has to be small since it has been shown experimentally that the interference landscape created by static WiFi interferers changes at these timescales when a mobile wireless sensor network moves at pedestrian speeds [3]. Hence, longer observation windows would not provide useful insights about the link quality.

The amount of observable quantities from an IEEE 802.15.4-compliant transceiver is limited to the Packet Reception Rate (PRR), the Received Signal Strength (RSS), and the Link Quality Indicator (LQI) associated with the received packet. Hitherto, there has been a long-standing debate in the research community on which of these indicators is the most appropriate for assessing the link quality properly, but no full agreement has been reached, especially when considering only a few packet samples [4]. This led to the exploration of hybrid metrics that make use also of the data coming from other layers in order to maximize the efficiency of data collection protocols [5], and to the creation of tools to analyze the statistical properties of link quality metrics in large static testbeds [6].

However, since our first concern is to minimize the size of the observation window and to obtain a fast assessment – ideally in the order of ten packets or one second of time – we need to deal with the following limitation: the fewer observed packets, the noisier become the RSSI, LQI, and PRR estimates. For this reason, in order to improve the characterization of a link and still achieve good predictions, we need a single metric that combines the strengths (and mitigates the weaknesses) of

all the available hardware indicators.

Resting on these ideas, we design the *Triangle Metric*, a geometrical combination of hardware indicators that conveys the strengths of the physical observables LQI, RSSI, and PRR, into a single metric. We show how the geometrical combination and the introduction of the *window mean* enables a good assessment of the link quality also when considering a limited amount of packet samples.

We evaluate the prediction performance of the triangle metric in comparison with the individual indicators using different data sets obtained from measurements in static and mobile environments. Our evaluation shows how the triangle metric can classify links even when only few sample packets are available in the observation window, a common situation in mobile wireless sensor networks.

The paper proceeds as follows. Section II provides an analysis of the limitations of the common hardware indicators. In Section III, we discuss the design of our triangle metric, and in Section IV we present the results of an experimental evaluation in static and mobile settings. After reviewing the related work in Section V, we offer our conclusions in Section VI.

II. LIMITATIONS OF HARDWARE METRICS

When using an IEEE 802.15.4-compliant transceiver [1], there are only a limited amount of quantities that can be directly observed and delivered to a prediction scheme: the Packet Reception Rate (PRR), the Received Signal Strength Indicator (RSSI), and the Link Quality Indicator (LQI) associated with the received packet. The RSSI value is an average of the received signal strength at the packet's arrival time. When sampling the RSSI value at a point in time in which there is no ongoing transmission (for example in the gap between data and acknowledgement packet, see e.g. [7]), one can also measure the RSSI noise floor, which gives a direct indication of the amount of interference, and by subtracting it to the RSSI values, one can assess the SNR. The LQI is implementation-dependent: in the case of the Chipcon CC2420 radio transceiver [2] it gives an indication of the chip error rate.

Based on those four basic observable quantities, the community has suggested several metrics for link estimation by considering the mean or the variance of PRR [8], [9], [10], LQI [11], [12], RSSI [4], [13] and to a lesser extent SNR.

In this work, we show that these approaches have three major limitations. First, for a small number of samples, the mean and variance are susceptible to noise, which leads to an unclear distinction of the link quality. The PRR-based approach requires a relatively large number of observations to obtain usable results, and the RSSI- and LQI-based approaches are only partially able to map accurately to PRR. Second, each metric provides a different type of information, and instead of leveraging on the combined knowledge, most of existing studies analyze these metrics individually or on a pairwise basis. Furthermore, the PRR and the observed LQI/RSSI statistics are often computed on the *same* set of packets, and no attempt is made to predict the *future* PRR.




Fig. 1. Correlation between Packet Reception Rate (PRR) and (a) the mean SNR, (b) the standard deviation of the SNR, (c) the mean LQI (c), and (d) the standard deviation of the LQI.

In the reminder of this section, we illustrate in greater detail the weaknesses and strengths of each basic indicator. Our aim is to highlight their differences, but also the mutual complementarity, which is the main motivation of our study.

Packet Reception Rate (PRR). The main limitation of the PRR is that it cannot differentiate between good stable links (i.e., links that provide a $PRR = 1$ and that are also resilient to external effects), and good links that might be unstable (i.e., links that can have a $PRR = 1$ but any minor environmental change such as shadowing or interference can significantly degrade their quality [4]). Furthermore, the smaller the dataset on which the PRR is computed, the lower the granularity. For example, when the PRR is computed on a set of only 10 packets, the reception of a single packet has an impact of 10% on the total result.

Signal to Noise Ratio (SNR). The signal-to-noise-ratio and the RSSI (Received Signal Strength Indicator) have been extensively studied in the literature [4], [11], [14]. In general, the characteristics of the SNR complement to some extent the limitations of the PRR: the latter cannot differentiate between good and very good links, but it can approximately differentiate between bad ($PRR < 0.35$), average ($0.35 \leq PRR \leq 0.75$) and good links ($PRR > 0.75$). On the other hand, SNR can *only* differentiate between very good links and the rest. As shown in Figure 1(a), a link with a mean SNR above 20 dB can be safely considered a very good link, but links with average SNRs between 5 dB and 10 dB are hardly distinguishable between bad, average and good. Hence, while PRR and SNR cannot accurately classify the entire spectrum of link qualities independently, combining their information could improve the classification process. Continuing with Figure 1(a), a link with $PRR > 0.75$ and a relatively high SNR, say 12 dB, can be easily identified as a good link, while the classification of a link with a mean SNR of 7.5 dB can be improved by using the PRR information.

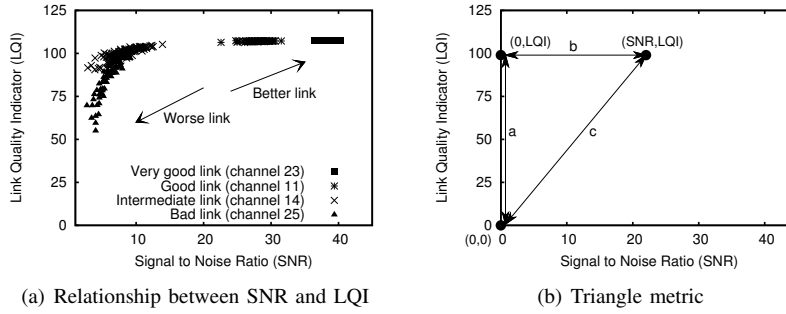


Fig. 2. The higher the SNR and the LQI, the better the link quality (a). This observation motivates the geometric basis of our approach: the link quality can be estimated by computing the distance of the point (SNR, LQI) from the origin $(0,0)$, i.e., by calculating the length c of the hypotenuse of the triangle abc (b).




Fig. 3. Examples of how to apply the triangle metric to the computation of the link quality. The length c of the hypotenuse, and thus the distance from the origin, increases as long as the link quality increases (a,b), while it decreases with lower-quality links (c,d) that cannot sustain a high packet reception rate.

Figure 1(b) shows the PRR over the standard deviation of the SNR. This combination provides negligible estimation information. The reason behind this is two-fold. First, the well-known log-normal shadowing model states that the SNR variance is constant for any mean SNR, henceforth, no links (good or bad) have a particular variance that differentiates them. Second, when few samples are provided, the variance is smaller for smaller sets than larger sets. Hence, a link with two received packets would have a misleading lower variance and hence would appear more stable than a link with 10 packets received.

The results shown in Figure 1 are based on data retrieved from 10 nodes deployed in a dynamic environment (cafeteria). Each node follows a TDMA scheme to broadcasts 10 continuous packets, one every second. The TDMA scheme is used to avoid internode interference, but other sources of interference are considered, such as 802.11b networks. Upon reception of a 10-packets batch, we calculate the PRR, and the mean and standard deviation of SNR and LQI.

The Link Quality Indicator (LQI). The characteristics of LQI are similar to those of PRR: like PRR, LQI presents a saturation that makes it incapable to distinguish between good and very good links. On the other hand, LQI shows a smoother decay that enables a better classification of bad, average, and good links [11]. Figure 1(c) shows the relationship between PRR and mean LQI. Despite the smoother decay of LQI, a link with only 2 received packets can still be confused with a link with 8 received packets, since they might have the same

LQI value. Hence, similarly to PRR and SNR, the LQI by itself is not sufficient to assess the quality of a link.

Some authors have suggested the use of the LQI variance for link classification [4], [12] based on two premises: (i) the better the link, the better the LQI, and (ii) the LQI reaches a saturation point (108 for the CC2420): hence, good links have a lower variance. This proposal holds when considering a moderate number of sample packets: when few samples are sent, this approach suffers from the intrinsic limitation of the variance explained before: for a small set of samples, fewer receptions lead to a lower variance, and hence, bad links could be classified as good, and vice versa. Figure 1(d) shows the low correlation between PRR and the LQI standard deviation.

III. THE TRIANGLE METRIC

The aim of our work is to combine the link state information of PRR, SNR, and LQI into a single metric that can accurately estimate the goodness of the link also with small observation windows. Given the resource limitations, we consider the mean values of PRR, SNR, and LQI as our input parameters. We do not consider the variance of those values because of the limited resources available on the sensor nodes, and for the problems occurring when considering only a small number of samples that we illustrated in Section II.

The mean SNR and the mean LQI provide important insights into the quality of a wireless link, as explained in the previous section. Based on the information provided by these two indicators, it is obvious that a good wireless link should have both a high mean SNR and a high mean LQI

at the same time. Following this observation, and illustrating it geometrically as depicted in Figure 2, we can introduce the key idea of our approach: the higher the SNR and the higher the LQI, the better the link (Figure 2(a)), hence, the link quality can be estimated by calculating the distance of the point $(\overline{SNR}, \overline{LQI})$ from the origin (0,0) (Figure 2(b)). The data plotted in Figure 2(a) is retrieved from a data collection application running on static nodes deployed in an office environment, communicating on different radio channels. We can see how the communication on different channels results in different link qualities, probably due to the effect of WiFi interference.

Figure 3 shows how the distance of the point $(\overline{SNR}, \overline{LQI})$ from the origin (0,0) can be used to estimate the quality of the link: different links will lead to different distances, and the better the link, the higher the distance from the origin.

Despite that the mean SNR and the mean LQI provide important insights about the quality of the link, they still carry a limitation: they only consider information from the received packets and disregard the information provided by lost packets. For example, a short-term effect may lead to the reception of only one packet with a high SNR and LQI on a link, while a different link may have received 8 packets with a mean SNR and mean LQI similar to the former. As a consequence of under-sampling, these two links could be erroneously classified as nearly equivalent. Figure 4(a) illustrates this problem: the blue circles represent $PRR = 1.0$, the black circles represent $PRR \in [0.75, 1.0)$, the green circles represent $PRR \in [0.35, 0.75)$, and the red circles represent $PRR \in (0, 0.35)$. The distance of the points $(\overline{SNR}, \overline{LQI})$ from the origin (0,0) do not correlate clearly with the packet reception rate. Hence, PRR still carries important link information not captured by mean SNR and mean LQI.

In order to include the PRR information, instead of using the statistical mean, we sum the SNR and LQI values of the received packets and divide it by the total number of transmitted packets. We call this operation the *window mean*. The advantage of the window mean over the statistical mean is that it includes the reception rate information by penalizing links with low reception rates. Figure 4(b) depicts the same links as in Figure 4(a), but using the window mean of SNR and LQI. As we can observe, the combination of PRR, SNR, and LQI leverages on the individual advantages of each metric and provides a more accurate differentiation of the link quality.

The formal description of our triangle metric is the following: let us assume that n packets are used to sample the channel and m of those packets are successfully received ($0 < m \leq n$). The LQI and SNR of each successfully received packet i are denoted by lqi_i and snr_i . Upon reception of the sampling packets, the receiver calculates the window mean SNR and LQI in the following way:

$$\overline{SNR}_w = \frac{\sum_{k=1}^m snr_k}{n} \quad (1)$$




Fig. 4. The combination of PRR, SNR, and LQI provides a more accurate differentiation of the link quality. There is no high correlation between the PRR and the distance of the point $(\overline{SNR}, \overline{LQI})$ from the origin (a). This implies that PRR still carries important link information not captured by the mean SNR and mean LQI, and including the PRR information using the window mean provides a more finegrained assessment of the link quality (b).

$$\overline{LQI}_w = \frac{\sum_{k=1}^m lqi_k}{n} \quad (2)$$

Then, the receiver calculates the distance to the origin (length of hypotenuse):

$$d_{\Delta} = \sqrt{\overline{SNR}_w^2 + \overline{LQI}_w^2} \quad (3)$$

Based on this distance, the receiver estimates the quality of the sender-receiver link according to a rule in which the larger the distance, the higher the link quality. In this work we assign empirical-based thresholds th to differentiate the quality of the links as follows:

$$\Gamma = \begin{cases} \text{Very Good link,} & th_{good} < d_{\Delta} \\ \text{Good link,} & th_{avg} \geq d_{\Delta} < th_{good} \\ \text{Average link,} & th_{bad} \geq d_{\Delta} < th_{avg} \\ \text{Bad link,} & d_{\Delta} < th_{bad} \end{cases} \quad (4)$$

IV. EVALUATION

We evaluate the triangle metric on static and mobile scenarios, and verify that in both cases it actually combines the strengths of its input metrics, outperforming them. We carry out the estimation using variable-sized observation windows, and we compare the results to the future packet delivery rate. Our goal is also to find a minimum window size that still gives reliable link estimations.

For the purpose of this paper, we need to create categories for link qualities, similar to the ones selected by Srinivasan et al. [15], because as pointed out in Section II, when using a small observation window, e.g., 10 packets, the impact of a single reception on the overall PRR is 10%, i.e., very high. We divide the links into four categories based on empirical observations, and we follow the distinction made in Section II. Therefore, we distinguish between very good, good, intermediate, and bad links. A very good link is intended to be a link with $PRR = 1$ as in [15], while a good link is a link with $PRR > 0.75$. An intermediate link has a packet reception rate between 0.35 and 0.75, so it is characterized by some

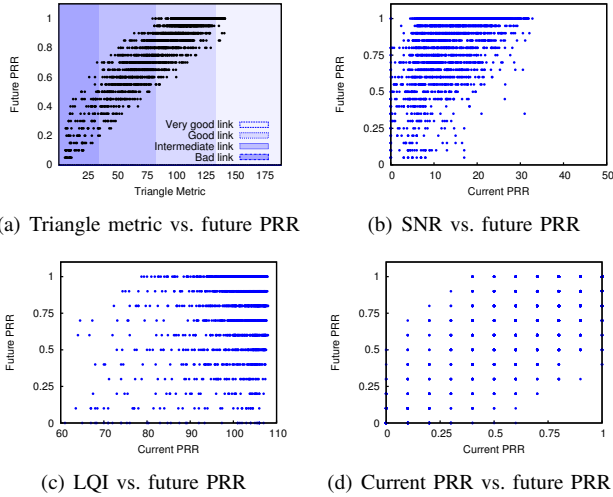


Fig. 5. Relationship between the link quality metrics and the future PRR when using an observation window of 20 packets for a mobile link.

persistent packet loss rate, while a bad link is a link whose packet reception rate is below 0.35.

For enabling a direct comparison, we also need to define a similar classification for the individual SNR and LQI, and for the triangle metric. We derive these thresholds from our investigation carried out in the previous sections. In particular, the thresholds of the triangle metric, SNR, and LQI can be mapped to the lines c, b, and a in Figure 3, respectively. Table I summarizes the categories and the selected thresholds.

TABLE I
CLASSIFICATION OF LINKS, AND THEIR RESPECTIVE THRESHOLDS.

Category \ Metric	PRR	SNR	LQI	Triangle
Very good link	1	30+	106+	145+
Good link	0.75 - 1	15 - 30	102 - 106	80 - 145
Intermediate link	0.35 - 0.75	5 - 15	80 - 102	30 - 80
Bad link	0 - 0.35	0 - 5	0 - 80	0 - 30

A. Experimental setup

To evaluate if the triangle metric actually combines the strengths of the single indicators, we carry out several experiments with real nodes. All nodes run the Contiki operating system [16] on Sentilla Tmote Sky [11] and Sentilla JCreate nodes. Both platforms are equipped with the 2.4 GHz Chipcon CC2420 radio transceiver. We evaluate the different metrics in a static and a mobile scenario.

In the static scenario, we run a simple data collection between several pair of nodes at various distances and using different transmission powers. Each receiving node collects the information about SNR, LQI, and sequence numbers of the received packets and logs them into a text file. We transmit data on all the channels of the 2.4 GHz band periodically with dedicated time-slots. All transmitters send 64 unicast packets per second. We extract the information available in the log files, select the size of the observation and prediction windows, and compute the amount of received packets, the mean SNR, and the mean LQI.

TABLE II
EVALUATION WITH STATIC LINKS

Assessment \ Future PRR	$x > 0.75$	$0.35 < x < 0.75$	$x < 0.35$
Metric: Triangle			
Very good link	100%	-	-
Good link	95%	5%	-
Intermediate link	36%	57%	7%
Bad link	3%	18%	79%

Metric: Current PRR			
Very good link	99%	1%	-
Good link	77%	22%	1%
Intermediate link	39%	56%	6%
Bad link	4%	23%	73%

Metric: Mean SNR			
Very good link	100%	-	-
Good link	94%	6%	-
Intermediate link	66%	29%	4%
Bad link	14%	33%	54%

Metric: Mean LQI			
Very good link	94%	6%	-
Good link	89%	10%	1%
Intermediate link	73%	23%	4%
Bad link	6%	16%	79%

As in the static scenario, also in the mobile scenario we have different pairs of nodes that periodically transmit packets at a rate of 64 packets per second. While the receivers have fixed positions, the transmitters are placed on humans moving at a constant pedestrian walking speed of approximately 5 km/h. The nodes move continuously in and out off the receiver's radio range.

B. Static scenario

Given thirty thousands transmissions from static links, we initially select an observation window of 20 packets and a prediction window of 250 packets. Figure 5 shows that given the limited size of the observation window, PRR, SNR, and LQI do not have a linear relationship with the future PRR, because they suffer from the limitations highlighted in Section II. On the contrary, the triangle metric has a more linear relationship, because it combines the information embedded in all the individual metrics.

We further carry a longer evaluation with observation and prediction windows of 10 packets, and the results are shown in Table II. As we can see, the mean SNR gives an acceptable estimation of good and very good links, but it fails to provide a satisfactory estimation for bad and especially intermediate links. A link estimated as intermediate by SNR in 69% of the cases has a future PRR of more than 0.75. The LQI is useful to estimate bad links but does not perform well for other kinds of links. The PRR is good for intermediate links, while it often misclassifies good links as intermediate. Our experimental results show that the triangle metric is able to merge the information in a way that leads to a more robust estimation compared to the input metrics taken individually: this is its main strength. The highlighted cells of Table II show how the values returned by the triangle are actually combining the strength of each individual metric.

TABLE III
EVALUATION WITH MOBILE LINKS

Assessment \ Future PRR	$x > 0.75$	$0.35 < x < 0.75$	$x < 0.35$
Metric: Triangle			
Very good link	99.5%	0.3%	0.2%
Good link	91%	9%	-
Intermediate link	41%	47%	12%
Bad link	8%	38%	54%
Metric: Current PRR			
Very good link	98%	2%	-
Good link	86.5%	13.3%	0.2%
Intermediate link	47%	43%	10%
Bad link	8%	45%	48%
Metric: Mean SNR			
Very good link	99.4%	0.4%	0.2%
Good link	93%	7%	-
Intermediate link	57%	36%	7%
Bad link	16%	50%	34%
Metric: Mean LQI			
Very good link	96%	3%	1%
Good link	71%	27%	2%
Intermediate link	45%	41%	14%
Bad link	-	36%	64%

The quality of the estimation provided by the triangle metric increases with the size of the observation window, since more data is available then. Figure IV-C shows the correlation between the size of the observation window and the future PRR (prediction window of 250 packets). The results in the figure show that a window size of 10 packets is a reasonable trade-off between accuracy, cost, and estimation time.

C. Mobile scenario

In the case of mobile sensor nodes, the link estimation is not as good as in the case of static scenarios. Nevertheless, our results shown in Table III confirm that the triangle metric combines the indicators so that the result is more accurate than the one provided by the individual indicators. The triangle metric estimates very good and good channels with comparable reliability w.r.t. SNR, intermediate links w.r.t. PRR and bad links w.r.t. LQI, as for static scenarios. This confirms that even in highly mobile scenarios, the triangle metric leads to a more robust assessment with respect to the input metrics taken individually.

It is important to remark that the results presented in Tables II and III depend on the values selected for the thresholds (Table I). We estimated these thresholds based on empirical observations, and hence, they may not be optimal. However, we took care in selecting thresholds that captures the common understanding that the community has of what represents very good, good, intermediate and bad links [15].

V. RELATED WORK

Link estimation gained significant attention in the WSN community after some initial works [17], [18], [19], [20] highlighted the unreliable nature of sensor network links. These studies showed that the transmission coverage of wireless sensor networks consists of 3 regions: (i) connected, where links are reliable and symmetric, (ii) disconnected, where there




Fig. 6. Impact of the size of the observation window: larger observation windows lead to more accurate estimations.

are no practical links for communications, and (iii) transitional, where a large percentage of links are unreliable. Since the extent of the transitional region can have a major impact on the performance and energy consumption of sensor nodes [21], [22], the research community found it necessary to investigate link estimators to improve the overall performance of the network.

Exploiting the metrics offered by IEEE 802.15.4-compliant transceivers, the community has mainly focused on four metrics: the packet reception rate (PRR), the signal-to-noise ratio (SNR), the received signal strength indicator (RSSI), and the link quality indicator (LQI).

Two well-known metrics based on the PRR are ETX [8], [9] and the requested number of transmissions [10]. However, as presented in Section II, PRR gives a limited estimation accuracy, since it cannot differentiate between stable and unstable good links, and its utility decreases when shortening the size of the observation window.

RSSI and LQI have been the core of a long dispute among researchers. Lin et al. [13] show experimentally that there are RSSI and LQI thresholds beyond which a link can sustain a PRR of at least 95%. However, for values below these thresholds, both metrics cannot be used to differentiate links clearly. Some studies suggested that the mean LQI is generally a better indicator because of the greater linearity with respect to the packet reception rate [11], [23], [24]. Other studies, instead, suggest that due to the high variance of LQI, at least 120 packets are needed to obtain a reliable estimation of average links when using the LQI, and that this drawback makes RSSI a better estimator than LQI [4], [25], [26].

Our work differs from the studies described above in two important ways. First, instead of focusing on a single metric, we combine the positive features of all of them into a more robust estimator. Secondly, instead of using the mean or the variance, we use the window mean, a different estimator that enables the reduction of the noise generated when evaluating a limited number of samples. Our work further shows that the requirement of using only few samples affects significantly the accuracy of variance estimators and hence they should not be considered.

A work closer in spirit to ours is the one presented in [14], where the authors propose a multiplicative metric between PRR and RSSI to estimate the link quality, but that does

not include the LQI, nor does it aim at fast estimation. Our goal is instead to design a suitable metric for mobile wireless sensor networks, i.e., one that returns a *fast* and *reliable* assessment also with small observation windows. We thus focus on fast estimation, and we do not attempt to estimate the burstiness as in [27], but the delivery capacity. Furthermore, differently from [5], we do not combine information from different layers, rather we exploit the information available in the radio transceiver, and we do not focus on ad-hoc wireless mesh but on mobile wireless sensor networks.

VI. CONCLUSIONS

The problem of link quality estimation has received significant attention from the wireless sensor network community due to its central role on the performance of two important network-wide metrics: delivery rate and energy consumption.

Our aim is to obtain a suitable link estimator for mobile wireless sensor networks, i.e., one that enables a fast assessment and minimizes the traffic overhead while still providing a reliable estimation. Towards this end, we propose the triangle metric, a link quality metric that combines geometrically the strengths of PRR, SNR, and LQI into a more robust estimator.

Our results show that the triangle metric provides a quick and reliable estimation with as few as 10 packets, and that it performs well both in static and dynamic environments.

ACKNOWLEDGMENTS

This work has been partially supported by the European Commission under FP7 with contract FP7-2007-2-224053 (CONET, the Cooperating Objects Network of Excellence). This research has been also partially financed by SSF, by VINNOVA, the Swedish Agency for Innovation Systems, by the IRCSET Postdoctoral fellow Grant PD200857, SFI Grant No. SFI08CEI1380, and by the Cluster of Excellence 306/1 "Inflammation at Interfaces" (Excellence Initiative, Germany).

REFERENCES

- [1] IEEE 802.15.4 Working Group. *Wireless Medium Access Control (MAC) and Physical Layer (PHY) specifications for Low-Rate Wireless Personal Area Networks (WPANs)*, September 2006.
- [2] Chipcon AS. *CC2420 RF Transceiver datasheet (Rev. B)*, March 2007.
- [3] J. Hauer, V. Handziski, and A. Wolisz. Experimental Study of the Impact of WLAN Interference on IEEE 802.15.4 Body Area Networks. In *Proc. of the 6th European Conference on Wireless Sensor Networks (EWSN)*, Cork, Ireland, February 2009.
- [4] K. Srinivasan and P. Levis. RSSI is Under Appreciated. In *Proc. of the 3rd Workshop on Embedded Networked Sensors (EmNets)*, Cambridge, MA, USA, May 2006.
- [5] R. Fonseca, O. Gnawali, K. Jamieson, and P. Levis. Four-Bit Wireless Link Estimation. In *Proc. of the 6th Workshop on Hot Topics in Networks (HotNets)*, Atlanta, Georgia, USA, November 2007.
- [6] N. Baccour, M. Ben Jamaa, D. do Rosario, A. Koubaa, M. Alves, L.B. Becker, H. Youssef, and H. Fotouhi. RadialLE: a Framework for Benchmarking Link Quality Estimators. In *Proc. of the 7th European Conference on Wireless Sensor Networks (EWSN), demo session*, Coimbra, Portugal, February 2010.
- [7] J. Polastre, J. Hill, and D. Culler. Versatile Low Power Media Access for Wireless Sensor Networks. In *Proc. of the 2nd Conference on Embedded Networked Sensor Systems (SenSys)*, Baltimore, MD, USA, November 2004.
- [8] D.S.J. De Couto, D. Aguayo, J. Bicket, and R. Morris. A High-throughput Path Metric for Multi-hop Wireless Routing. In *Proc. of the 9th Conference on Mobile Computing and Networking (Mobicom)*, pages 134–146, San Diego, CA, USA, September 2003.
- [9] D.S.J.D. Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. *Wireless Networks*, 11(4):419–434, 2005.
- [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. of the 6th Symposium on Mobile ad hoc networking and computing (MobiHoc)*, pages 414–425, Urbana-Champaign, IL, USA, May 2005.
- [11] J. Polastre, R. Szewczyk, and D. Culler. Telos: Enabling Ultra-Low Power Wireless Research. In *Proc. of the 4th Symposium on Information Processing in Sensor Networks (IPSN)*, Los Angeles, USA, April 2005.
- [12] C.A. Boano, T. Voigt, A. Dunkels, F. Österlind, N. Tsiftes, L. Mottola, and P. Suárez. Exploiting the LQI Variance for Rapid Channel Quality Assessment. In *Proc. of the 8th Conference on Information Processing in Sensor Networks (IPSN), poster session*, pages 369–370, San Francisco, CA, USA, April 2009.
- [13] S. Lin, J. Zhang, G. Zhou, L. Gu, T. He, and J.A. Stankovic. ATPC: Adaptive Transmission Power Control for Wireless Sensor Networks. In *Proc. of the 4th Conference on Embedded Networked Sensor Systems (SenSys)*, Boulder, Colorado, USA, November 2006.
- [14] M. Rondinone, J. Ansari, J. Riihijärvi, and P. Mähönen. Designing a reliable and stable link quality metric for Wireless Sensor Networks. In *Proc. of the Workshop on Real-world Wireless Sensor Networks (RealWSN)*, pages 6–10, Glasgow, Scotland, April 2008.
- [15] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis. An Empirical Study of Low Power Wireless. In *SING Tech Report*, October 2008.
- [16] A. Dunkels, B. Grönvall, and T. Voigt. Contiki - a Lightweight and Flexible Operating System for Tiny Networked Sensors. In *Workshop on Embedded Networked Sensors*, Tampa, Florida, USA, November 2004.
- [17] J. Zhao and R. Govindan. Understanding Packet Delivery Performance in dense Wireless Sensor Networks. In *Proc. of the 1st Conference on Embedded Networked Sensor Systems (SenSys)*, pages 1–13, Los Angeles, CA, USA, November 2003.
- [18] A. Woo, T. Tong, and D. Culler. Taming the Underlying Challenges of Reliable Multihop Routing in Sensor Networks. In *Proc. of the 1st Conference on Embedded Networked Sensor Systems (SenSys)*, pages 14–27, Los Angeles, CA, USA, November 2003.
- [19] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker. An Empirical Study of Epidemic Algorithms in Large Scale Multihop Wireless Networks. *Technical Report IRB-TR-02-003*, March 2002.
- [20] A. Cerpa, N. Busek, and D. Estrin. Scale: A tool for simple connectivity assessment in lossy environments. *CENS Technical Report 0021*, September 2003.
- [21] G. Zhou, T. He, S. Krishnamurthy, and J.A. Stankovic. Impact of Radio Irregularity on Wireless Sensor Networks. In *Proc. of the 2nd Conference on Mobile systems, applications, and services (MobiSys)*, pages 125–138, Boston, MA, USA, June 2004.
- [22] M.A. Zuniga, K. Seada, B. Krishnamachari, and A. Helmy. Efficient geographic routing over lossy links in wireless sensor networks. *ACM Transactions in Sensor Networks*, 4(3):1–33, 2008.
- [23] M.M. Holland, R.G. Aures, and W.B. Heinzelman. Experimental Investigation of Radio performance in Wireless Sensor Networks. In *Proc.s of the 2nd Workshop on Wireless Mesh Networks (WiMesh)*, Reston, Virginia, USA, September 2006.
- [24] T. Rein. Energy and Time Efficient Link-Quality Estimation for Wireless Sensor Networks. Master's thesis, ETH Zurich, April 2007.
- [25] S. Xiao, A. Dhamdhare, V. Sivaraman, and A. Burdett. Transmission Power Control in Body Area Sensor Networks for Healthcare Monitoring. In *IEEE Journal on Selected Areas in Communications*, volume 27, pages 37–48, January 2009.
- [26] A. Meier, T. Rein, J. Beutel, and L. Thiele. Coping with Unreliable Channels: Efficient Link Estimation for Low-Power Wireless Sensor Networks. In *Proc. of the 5th Conference on Networked Sensing Systems (INSS)*, Kanazawa, Japan, June 2008.
- [27] K. Srinivasan, M. Kazandjieva, S. Agarwal, and P. Levis. The Beta Factor: Measuring Wireless Link Burstiness. In *Proc. of the 6th Conference on Networked Embedded Sensor Systems (SenSys)*, Raleigh, NC, USA, November 2008.