

Prediction Accuracy of Link-Quality Estimators

Christian Renner, Sebastian Ernst, Christoph Weyer, and Volker Turau

Hamburg University of Technology, Hamburg, Germany,
{christian.renner,c.weyer,turau}@tu-harburg.de

Abstract. The accuracy of link-quality estimators (LQE) is mission-critical in many application scenarios in wireless sensor networks (WSN), since the link-quality metric is used for routing decisions or neighborhood formation. Link-quality estimation must offer validity for different timescales. Existing LQEs describe and approximate the current quality in a single value only. This method leads to a limited accuracy and expressiveness about the presumed future behavior of a link. The LQE developed in this paper incorporates four quality metrics that give a holistic assessment of the link and its dynamic behavior; therefore, this research is an important step to achieving a higher prediction accuracy including knowledge about the short- and long-term behavior.

1 Introduction

For most algorithms in wireless sensor networks (WSN) it is essential that each node has thorough knowledge about its direct neighbors. This information is collected and provided by neighborhood management protocols and is used, e.g., for routing decisions, group formation, or data sharing. The dynamic behavior over time of the wireless channel and the missing correlation between adjacency and possibility of communication—due to obstacles and multi-path propagation—render the definition of the neighborhood of a node a nontrivial task. One important criterion used by neighborhood management protocols to determine the importance of a node is the quality of the communication between nodes, which in turn is provided by a link-quality estimator (LQE). Depending on the intended application, either the short- or long-term qualities of a link or a combination of both is preferred for choosing an appropriate node in the vicinity [2]. In the past many different approaches were investigated [14,6,4,7,2,11,3,1] or are currently used, e.g., in TinyOS [10,8].

In principle a LQE measures the quality based on logical (e.g., packet success rate) or physical (e.g., received signal strength) metrics. Newer proposals use a combination of these in order to improve the accuracy of the prediction [1]. However, common to these LQEs is that the measured quality is squeezed into a single value—e.g., a moving average—due to memory restrictions and for easier

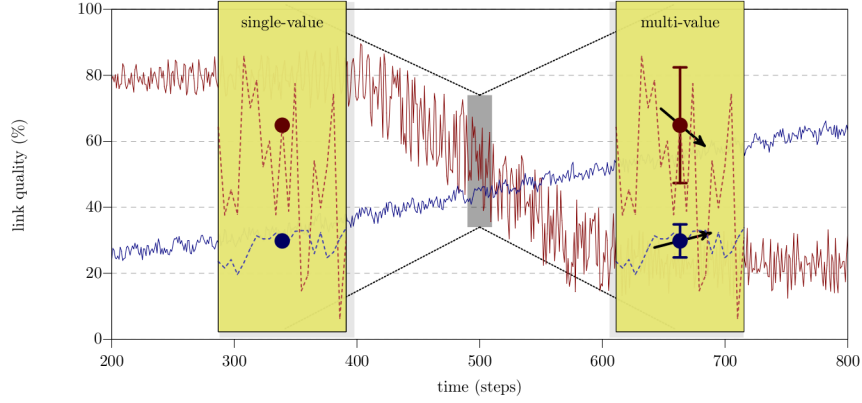


Fig. 1. Link-quality estimation of a deteriorating and an improving link with a single-value (left) and statistical multi-value (right) technique

comparability. In doing so, the value represents only a snap-shot of the plain link-quality at a specific point in time without any additional information about variation and the current trend of the long- or short-term behavior in the past. The expressiveness of such a single-value metric is limited: This procedure is comparable to a stock market, where a stock is described only by its current value or an average of its last values. Based on this limited information it is a game of luck to decide in which stock to invest in the future. Furthermore, it is important to know which stock will perform better for a short-term profit and which is suitable for a long-term investment.

The same considerations are valid for judging the quality of a link. Figure 1 visualizes the problem of comparing two link-qualities. With a metric that reflects only the current quality, a comparison is nearly impossible even for a short-term projection. Based on a single-value prediction the deteriorating link appears superior, since its current link-quality is higher. When using multi-value prediction the dynamic behavior of links can also be incorporated into the decision, so that the improving link will be favored: Although its current link-quality is lower, its positive trend and low deviation lead to a better future perspective when compared to the large variation and negative trend of the deteriorating link. Furthermore, all existing LQEs are only capable of providing either a short- or a long-term prediction. Yet, an application often needs different prediction windows, e.g., a long-term prediction for forming a cluster and a short-term one for selecting the eligible next routing hop.

In this paper we present a LQE that effectively calculates for each link four quality characteristics: short- and long-term quality, variation, and an indicator of the current trend. This LQE is a general approach that can be easily adopted to various sources of link-quality metrics. Intensive simulations and experiments are undertaken in order to proof the advantages of the newly developed approach over existing ones. We compare and analyze the ability of link-quality tracking

of modern and well-known LQEs with our new approach. In this context, we identify problems and benefits of the different estimators, and we also determine the correlation and average error with respect to empirical ground-truth measures. Moreover, we analyze the ability of the LQEs to select reliable links, which is commonly needed for routing decisions and cluster formation. The paper concludes with a discussion of the results and intended future research activities.

2 Link-Quality Estimation Techniques

Estimating link-qualities in WSNs is a nontrivial task. Many existing solutions for other wireless networks are not feasible for WSNs. Particularly, the lack of infra-structure, limited memory, energy constraints, and low-cost transceivers complicate possible solutions. Various approaches were proposed in the past to overcome these limitations and to provide a meaningful metric describing the actual link-quality and thus predicting its future behavior. In this section, a brief overview and discussion of general link-quality metrics is carried out, followed by a thorough view on already existing approaches employing them.

2.1 Link-Quality Metrics

Basically, there exist two different categories of link-quality metrics: physical and logical indicators. The former are provided by the radio hardware and are based on the signal strength of a received packet, such as the Received Signal Strength Indication (RSSI), the Link-Quality Indication (LQI), or Signal-to-Noise Ratio (SNR). The logical indicators estimate the link quality by keeping track of message losses. Examples of such metrics are: Packet Success Rate (PSR), Required Number of Packets (RNP) [4], or Expected Transmission Count (ETX) [6] for describing the effort needed to successfully transmit a packet.

Physical Metrics Using physical metrics has several advantages. The metrics come without any additional costs, since the measurement is performed by the receiver hardware every time a packet is received, or can at least be calculated easily in the case of SNR. Also only a small number of samples is needed to get a first approximation of the link-quality. Additionally, the metrics can be measured by utilizing any traffic on the wireless channel without the need of periodical broadcasts if the application produces enough traffic.

However, several research activities on this topic have shown that the physical metrics are of shortened use [12]. First of all, the metrics are strongly dependent on the receiver hardware—e.g., in our experiments with the Atmel RF230 many links have a RSSI value close to the sensitivity of the transceiver, but provide reliable packet reception at the same time. Secondly, the RSSI and LQI are only available for successfully received packets, but not in the case of packet loss. There exist different attempts to improve the quality of the measurements—e.g., RSSI calibration [5]—but these introduce a higher computation-complexity not suitable for WSNs. Another disadvantage is the expressiveness of these metrics

with respect to the application-related perception, i.e., the expected PSR. In [9] Lal et al. show that physical metrics cannot always predict the PSR, especially in case of long-range links with an RSSI near the sensitivity threshold of the radio chip. In those cases the SNR-to-PSR relation is not deducible at all.

Logical Metrics The advantages of logical metrics are that they do not depend on specific hardware characteristics and correlate directly with an application point of view, i.e., the ratio of successfully transmitted packets. This however leads to the problem that a node needs to track the ratio between the number of successfully received packets versus packet loss in an efficient way. A window-based approach consumes too much memory, and counting received and lost packets does not incorporate that recent events should be weighted higher. Hence, most of the logical metrics are calculated with an Exponentially Weighted Moving Average (EWMA) [14]. In contrast to physical metrics, LQEs based on this method rely on frequent packet transmissions in order to keep the link-quality estimates up-to-date.

Periodic broadcast packets are often used to achieve this goal and are also the most commonly brought-up drawback of these approaches, since they waste energy and occupy the wireless channel. However, the measurements—even of the physical link-quality estimators—are conducted by the receiver, but the information must be available at the sender. To perform the necessary exchange of information for each neighboring node, piggy-backing this data on top of application packets is risky due to increased packet size, especially in dense networks. Moreover, the time of information exchange depends on the network traffic of each node, so that large delays may occur and nodes make decisions based on outdated link-quality data. Thus, a periodic information exchange using dedicated broadcasts is always necessary when link-quality estimation is needed. This mitigates animadversion on the periodical broadcasts.

2.2 Link-Quality Estimators for Sensor Networks

Woo et al. define the Window Mean with EWMA (WMEWMA) [14], one of the first LQEs for estimating the PSR. In a previous work they also investigate and compare WMEWMA with different existing approaches. Thereafter, EWMA-based estimators have been widely adopted in WSNs.

ETX introduced by De Couto et al. [6] tries to estimate the number of transmissions that are necessary to send a packet successfully. The number of received packets within a fixed window is counted and compared to the number of expected packets that are periodically broadcasted by each node. The disadvantage of ETX is that it is only updated at the end of each window. A short window thus leads to a high fluctuation of the ETX metric and a long window to an infrequently updated ETX. Cerpa et al. [4] introduce RNP that incorporates the distribution of losses within the window. They observed that a link with consecutive losses should be rated lower than links with discrete losses. Four-Bit (4B) [7] is based on ETX with several enhancements. They use an EWMA

for estimating the ETX and a second one for smoothing the final 4B metric. Additionally, 4B uses additional information from the link and network layer.

The Link Estimation Exchange Protocol (LEEP) is based on 4B and part of the current TinyOS version [8]. LEEP constitutes a layer between the MAC protocol and the application or routing engine, respectively. Whenever a packet is sent, LEEP attaches additional information to that packet: a sequence number, a counter for the number of known neighbors, and the ID and in-bound link quality of these neighbors. If the space (left by the upper layer) in the packet is too small, a round-robin procedure is used. The main disadvantage of LEEP is that in case of consecutive packet losses the link-quality is not updated. In addition to the limited number of possible links, due to memory restrictions, this behavior can result in a full neighborhood table with nodes that no longer exist.

The Adaptive Link Estimator (ALE), introduced by Weyer et al. [13], is an EWMA filter of the measured PSR. Each node sends a beacon per time interval (a so-called round) and records received beacons from other nodes within this interval. At the end of each round, link qualities are recomputed. Upon reception of a beacon from a previously unknown node the associated link-quality is initialized with a value of 50% to achieve a shorter rise-time for new links. The weight of the EWMA is adapted to the quality of the link. For good links ALE uses a higher weight for more stable estimation, while links with a lower quality are estimated in an agile fashion for a faster reaction.

An approach using Kalman Filter Based Link-Quality Estimation is proposed by Senel et al. [11]. They filter the RSSI of successfully received packets with a Kalman filter and subtract the noise floor to obtain an estimation of the SNR. Finally, a PSR is derived by applying a hardware specific SNR-PSR mapping for the transceiver. This approach is very complex and suffers from the restricted correlation between SNR and PSR [9].

The Fuzzy Link-Quality Estimator (F-LQE) is proposed by Baccour et al. [1]. The idea is to combine four different metrics into a single quality indicator using Fuzzy Logic. These four metrics are SNR, PSR, link asymmetry level (ASL), and stability (SF). The SNR is calculated after each successful packet reception by subtracting the power on the channel directly after packet reception from the signal strength of the packet. An EWMA filter is used to obtain the PSR. From the latter, the asymmetry level is obtained by calculating the absolute difference between the unidirectional PSR values of a node pair. The stability is defined as the coefficient-of-variation of the most recent 30 PSR values of a link.

3 Holistic Packet Statistics

All of the approaches presented in Sect. 2.2 have in common that they squeeze link quality into a single value. We argue that there should be a multi-faceted representation of link quality. Hence, we have devised the concept of Holistic Packet Statistics (HoPS). HoPS is tailored to provide detailed information about the static and dynamic behavior of a link using four distinct descriptors of the link quality. Due to the troublesome nature of the expressiveness of SNR values

and the corresponding PSR matches—as laid out in Sect. 2.1—we take a different course than recent approaches and focus on the development of a more sophisticated link-quality assessment using logical metrics.

HoPS allows for utilizing enriched link information by granting access to all four link-quality metrics rather than hiding information by compressing them into a single value. Yet, this does not imply that there is no way to combine the values of HoPS into a single link-quality estimate. In this section, we define the four quality metrics of HoPS and present two possible solutions for dynamic link-quality assessment.

3.1 Link-Quality Descriptors

The four link-quality descriptors used by HoPS are measured at the receiver side by monitoring the packet success rate. Thus, a packet should contain a sequence number for the detection of packet loss. Additionally, the metrics should be updated frequently in a constant interval. Many presented LQE implementations (cf. Sect. 2.2) update only after a packet is successfully received; this leads to a wrong estimation if no more packets are received from a node.

Short-term Estimation is realized using a first-order EWMA filter for retrieving the in-bound PSR of a link in the recent past:

$$h_{\tau}^{\text{ST}} = \alpha \cdot h_{\tau-1}^{\text{ST}} + (1 - \alpha) \cdot q_{\tau} . \quad (1)$$

The choice of the coefficient α influences the sensitivity to short-term changes in link quality. In the straight-forward case, q_{τ} is a binary value indicating whether an expected packet at time τ was received. Another option would be using a windowed mean, if more than one packet is received from the same node within an update interval. This technique is comparable to WMEWMA.

Long-term Estimation is obtained by a second-order EWMA filter, i.e., the values of the first-order estimation in (1) are smoothed by

$$h_{\tau}^{\text{LT}} = \beta \cdot h_{\tau-1}^{\text{LT}} + (1 - \beta) \cdot h_{\tau}^{\text{ST}} . \quad (2)$$

To achieve a strong smoothing effect, choosing β larger than α is advisable.

The dynamics of a link are obtained by means of the lower and upper deviation of the short-term estimation from the long-term estimation. The sum of lower and upper deviation yields the average absolute deviation. Due to space constraints the detailed derivation is omitted. To track the time-variant changes of these deviations and to preserve memory, another EWMA filter is applied:

$$\delta_{\tau}^{+} = \gamma \cdot \delta_{\tau-1}^{+} + (1 - \gamma) \cdot \varphi(h_{\tau}^{\text{ST}}, h_{\tau}^{\text{LT}}) , \quad (3)$$

$$\delta_{\tau}^{-} = \gamma \cdot \delta_{\tau-1}^{-} + (1 - \gamma) \cdot \varphi(h_{\tau}^{\text{LT}}, h_{\tau}^{\text{ST}}) , \quad (4)$$

$$\text{with } \varphi(x, y) = \begin{cases} x - y, & \text{if } x > y \\ 0 & \text{else} \end{cases} . \quad (5)$$

These two values are utilized as follows.

Absolute Deviation Estimation is the estimated average absolute deviation

$$h_{\tau}^{\sigma} = \delta_{\tau}^{+} + \delta_{\tau}^{-}. \quad (6)$$

This value gives an impression of the stability of the link in terms of the variation of h_{τ}^{ST} around its mean h_{τ}^{LT} . Its calculation demands less computing power than the recursively determined standard deviation, while not being less expressive.

Trend Estimation is intended to reflect the course that a link is taking:

$$h_{\tau}^{\theta} = \delta_{\tau}^{+} - \delta_{\tau}^{-}. \quad (7)$$

A floating link—i.e., its long-term quality has had no notable changes in the past and is therefore not likely to change in the future—can be identified by values of h^{θ} close to zero. In contrast, positive values indicate an improving link, whereas negative values expose a deteriorating link. The absolute value of the trend indicates the slope of the current trend.

3.2 Theory in Praxis

At this point, we give a brief introduction to the interpretation and usage of the four link-quality metrics of HoPS. There are many possible solutions for the utilization of the four link-quality indicators of HoPS. For instance, a routing protocol could favor a less varying link over one with the same long-term quality. Another possibility would be to use the variation and trend in order to compute a lower bound link-quality with a given confidence. Hence, there is no silver bullet for an enriched employment of these values. However, we want to present two examples of merging the four HoPS ingredients into a prediction value.

The first approach is a dynamically adjusted link-quality estimator:

$$H_{\tau}^{\text{dyn}} = h_{\tau}^{\text{LT}} + \frac{|h_{\tau}^{\theta}|}{h_{\tau}^{\sigma}} \cdot (h_{\tau}^{\text{ST}} - h_{\tau}^{\text{LT}}) \quad \left(0 \leq \frac{|h_{\tau}^{\theta}|}{h_{\tau}^{\sigma}} \leq 1 \right) \quad (8)$$

This estimator describes a floating link ($h^{\theta} \approx 0$) by its long-term behavior and a massively changing link ($|h^{\theta}| \approx h^{\sigma}$) by its short-term estimate. Intermediate assessment is achieved using the relative behavior of trend and variation.

Approach number two is a confident long-term predictor:

$$H_{\tau}^{\text{pred}} = \begin{cases} h_{\tau}^{\text{LT}} + h^{\theta} - \omega \cdot h^{\sigma}, & \text{if } h^{\theta} \geq \omega \cdot h^{\sigma} \quad (0 \leq \omega < 1) \\ h_{\tau}^{\text{LT}} + h^{\theta} + \omega \cdot h^{\sigma}, & \text{if } h^{\theta} \leq -\omega \cdot h^{\sigma} \\ h_{\tau}^{\text{LT}} & \text{else} \end{cases} \quad (9)$$

In contrast to the previous method, a link is classified by the difference between trend and variation. If a clear trend can be identified—i.e., a relative threshold ω is exceeded—the long-term value is shifted correspondingly. The goal is to improve the prediction of the link behavior by incorporating the trend.

4 Evaluation Methodology

To evaluate and compare HoPS with the LQEs introduced in Sect. 2.2, we conducted simulations based on a 13-day real-world experiment. The purpose of this experiment is to gather real-world data in terms of the PSR and channel information, which is used by the TinyOS Simulator (TOSSIM) to feed the different LQEs. The main benefits of this approach are (i) detailed information about node connectivity in a real-world scenario, (ii) an identical runtime situation for all experiments, and (iii) repeatable experiments under real-world conditions.

4.1 Data Basis

A sensor-node testbed, consisting of 15 IRIS nodes with distances between 1 m and 40 m, was employed in our University building. All nodes were USB-powered. Every node broadcasted beacon messages within a fixed interval of 4 s with a random jitter using a uniform distribution. The transmit power was 0 dBm on radio channel 22. Packets consisted of a virtual payload of 13 bytes—containing a beacon sequence number—and a 13 byte MAC header. Clear Channel Assessment (CCA) was performed with at most 5 retries. We have not deactivated this feature, as we are interested in realistic connectivity behavior of a sensor network; moreover, a failed CCA implies that the sender experiences a high radio signal, but it does not necessarily imply the inability of the receiver to correctly receive a packet. Upon reception of a beacon, the following data is logged on the connected computer: node ID of sender and receiver, the sequence number, LQI, and SNR. From this data, a detailed track of packet receptions and corresponding physical channel quality was produced. More than 280 000 packets were transmitted per node.

4.2 Methodology and Metrics

We adjusted the physical layer of TOSSIM to reproduce our office environment. This layer utilizes the data basis to decide if a packet sent at a given time is received by other nodes in the network. When a node puts a radio packet onto the channel, this physical layer checks for all other nodes in the network, if they have received the experiment beacon in that time slot. The packet is only delivered to the nodes, for which this is the case. The result is a realistic, repeatable, and fair simulation environment.

We obtained the PSR ground truth by applying Hamming windows of lengths from 30 s to 60 min. The different LQEs are compared in terms of their ability to correctly represent the current and future course of each link. For this comparison, we used a normed cross correlation function (CCF) and the mean absolute error (MAE). Quality values are in the range of 0 to 100% with an 8-bit resolution; EETX values are converted to PSR. In case of F-LQE, the raw value is used. The implementations of LEEP and RNP have a resolution of tenth. ALE and HoPS run with 16-bit integers. To achieve fair conditions, we also ran an adapted version of LEEP using double precision.

4.3 Parameters

The parameters of the LQEs have been chosen as proposed in the corresponding papers. LEEP and RNP use an EWMA coefficient of 0.9 and update the estimations every 3 packets. F-LQE uses the fuzzy membership functions from [1] and a fuzzy parameter of 0.6. ALE uses filter coefficients of 0.9 in agile and 0.987 in stable link state, where PSR estimates of 86% and 74% serve as the thresholds to switch states. The parameters for HoPS are $\alpha = 0.9$ for short-term and $\beta = \gamma = 0.997$ for long-term and deviation estimation. HoPS initializes $h^{ST} = h^{LT} = 50\%$ for new links. The threshold for the stable long-term prediction is $\omega = 0.25$. These parameters have been determined using a Java GUI that was developed for detailed link assessment and link-quality filter design.

5 Evaluation Results

In this section, the most important evaluation results are presented. First, the ability of different LQEs of link-quality tracking is analyzed and compared using the testbed data introduced in Sect. 4.1. Second, their performance in terms of correct link selection is evaluated.

5.1 Link-Quality Tracking

The estimation methods of the different LQEs vary largely, so that light has to be shed on their ability of link-quality tracking. In case of good links with a stable PSR above 95% and a SNR well above 5 dB, all LQEs achieve good estimates. If, in contrast, links have a dynamic behavior, this picture changes.

Medium Link A medium link with a long-term PSR above 75%, notable short-term deviations, and SNR values below 5 dB is shown in Fig. 2. The second plot reveals that ALE tracks the long-term quality very well, but immediately toggles to the agile state, when its estimation falls below the 74% PSR threshold, and returns to the stable state after reaching the 86% PSR threshold. In the worst case, this may lead to oscillation between these states. F-LQE assigns a rather low value to the link, which is mainly caused by the low SNR values: the quality value is only about half the value of the long-term PSR. Instead of following the application-relevant PSR value, F-LQE mirrors the course of the SNR. This leads to the strong lag regarding the recovery phase towards the end of the displayed time window, where the short-term PSR is approaching the long-term measure even before the SNR value is increasing again. LEEP and RNP natively describe link-quality as EETX values with a resolution of tenth, so that their quality curve degenerates to a step function. RNP tracks variations of the link quality very quickly, but produces many overshoots. In contrast to this, LEEP offers a strong low-pass characteristic, but produces peak values. They occur on some rapid short-term PSR improvements. This behavior is alleviated, if the

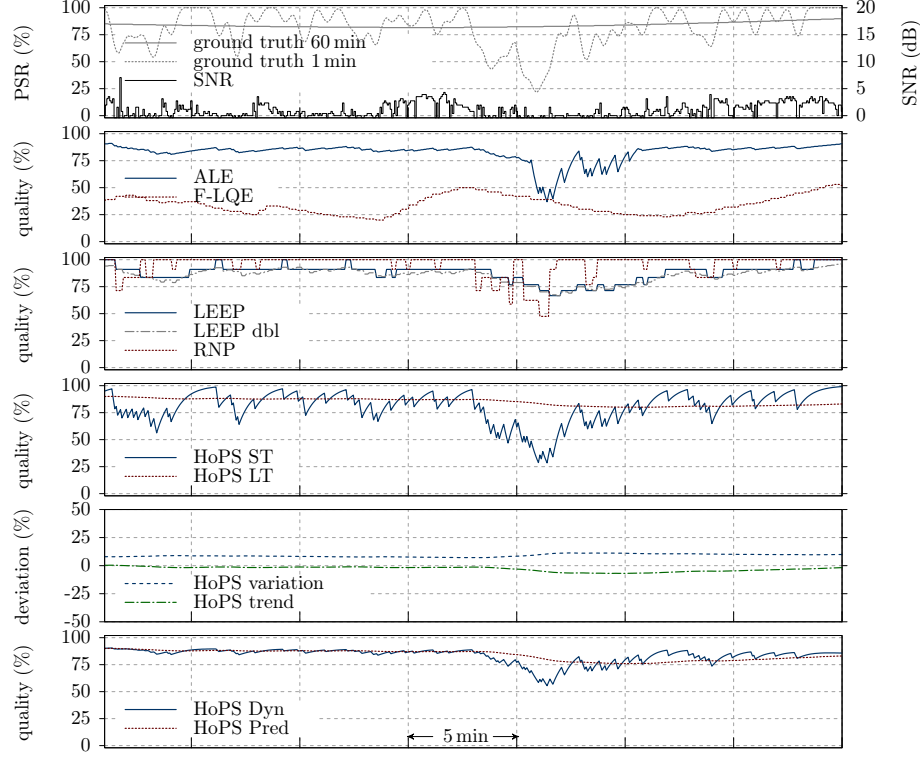


Fig. 2. LQE performance for a medium link with moderate to low SNR

resolution of LEEP is increased (LEEP dbl), so that LEEP can track link quality more smoothly.

The short- and long-term estimations of HoPS follow the corresponding ground-truth trace accurately, giving a differentiated picture of the link's course. In addition, the variation value produced by HoPS resembles the absolute average deviation of the two estimates in the recent past. The latter corresponds to the variation of the ground truth in the top plot. Up to the middle of the time window, the HoPS trend correctly identifies the link to be floating. With a small time delay, the decreasing short-term link-quality is identified. Due to the rather small and slow change, trend and variation stay below an absolute value of 10%. After the short disturbance of the link, the trend approaches zero and the variation is decreasing towards its old value. HoPS Dyn is in a stable long-term state until the beginning of the disturbance. Due to the increasing absolute value of the trend at that point, HoPS Dyn takes an intermediate value between short- and long-term prediction of HoPS. During the link recovery phase, HoPS Dyn slowly returns to the stable prediction. In contrast, HoPS Pred is only marginally influenced by the disturbance, since it is too short and small.

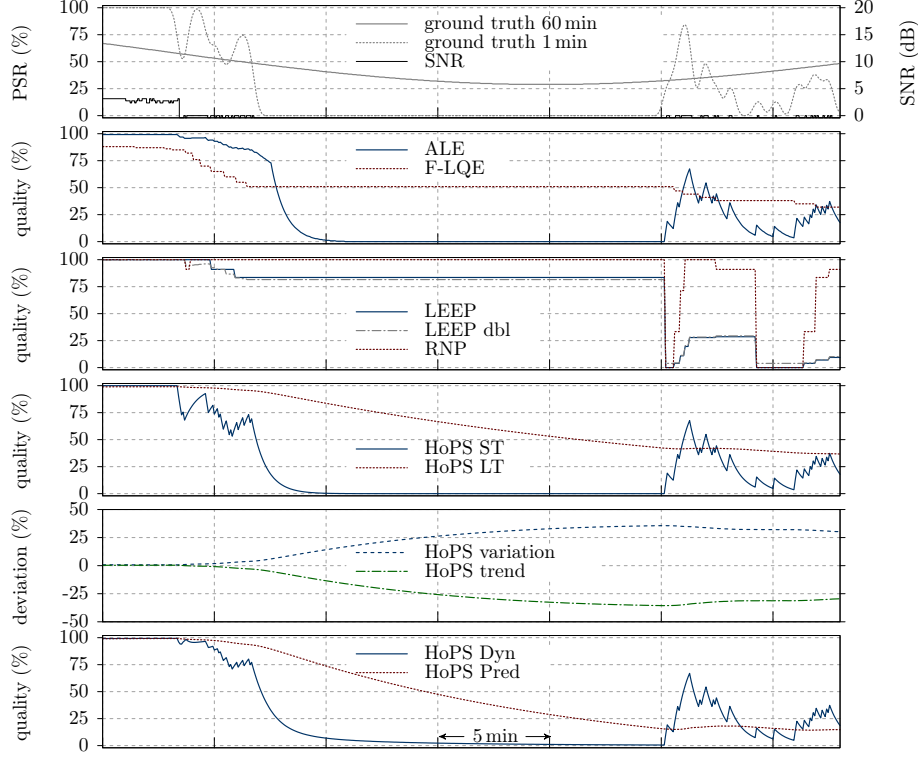


Fig. 3. LQE performance for a deteriorating link with low SNR

Poor Link A study of a deteriorating link reveals additional differences between the LQEs. The link displayed in Fig. 3 drops from a high, stable quality state to a completely useless link within less than 5 minutes.

ALE implicitly trusts the link for a few minutes due to the large filter coefficient. After falling below the 74% border line, ALE follows the quality drop quickly, but with a notable time-delay. The better a link has been, the larger this delay will be. Due to the adapted filter coefficient, the PSR bursts at the end of the displayed window are tracked closely. Although link tracking appears to work quite well, there are drawbacks. Firstly, there is no option to choose between short- and long-term estimation values. Secondly, random effects caused by the static thresholds for changing the link state may occur. The large peak after the long link disturbance misses the 86% threshold for changing to the stable state by less than 20%, which corresponds to the reception of only 2 packets.

F-LQE identifies the degenerating quality at the beginning accurately, since it coincides with a sudden SNR drop, but stops updating the quality estimate, because no more packets are received. Only upon the reception of new packets, the quality estimate is updated—which is happening quaintly when the link comes up again. Note that F-LQE assigns a larger link-quality value to the link

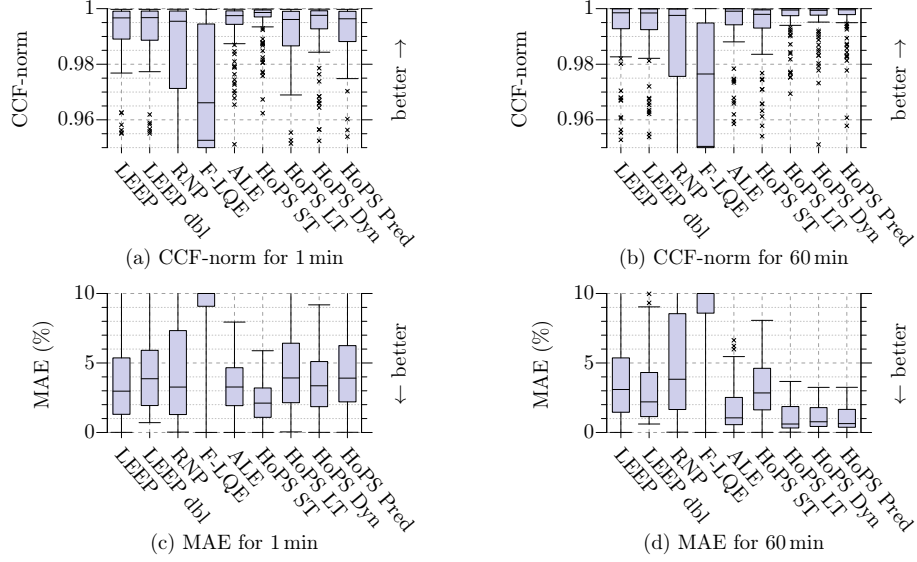


Fig. 4. Ability of the different LQEs to track the course of the empirical PSR

in its broken state than to the medium link in the previous example. RNP and LEEP reveal the same weakness, plus RNP overestimates the link during its bursty phase at the end. For an application, the broken link would look like a perfect (RNP) or good (LEEP) link for more than 15 minutes.

In contrast, HoPS keeps track of the declining link-quality for both the short- and long-term values. Due to the large filter coefficient of HoPS LT, there is a notable lag. The question on which value to trust is answered by the variation and trend: They take large values that have about the same absolute value. This indicates a massive change of the link away from its long-term behavior. Due to the small slope of HoPS LT, trend and variation keep their large absolute values until the end of the disturbance. At this point they are decreasing again, since the link is revitalizing. HoPS Dyn detects the deterioration phase quickly and switches to the short-term estimate. The previous dip in link-quality is ignored due to the almost zero trend at this point. During the bursty phase at the end, HoPS Dyn sticks to the short-term estimate, thus reflecting the uncertain future of the link. Note that this behavior can be adjusted by the choice of the deviation filter coefficient γ . HoPS Pred reacts slowly to the broken state of the link, but has a similar course as the long-term ground truth in the first plot. However, it does not react to the improving link-quality in a suitable way.

Statistical Analysis While these examples give a detailed understanding of the strengths and weaknesses of the LQEs, a statistical evaluation of all links is shown in Fig. 4. The box-and-whisker plots summarize the ability of link-quality tracking using the CCF-norm and the MAE (in absolute percent error) for all

links in the testbed, on which at least one packet was received. For short-term estimation ALE, HoPS ST and HoPS Dyn give the best results, see Fig. 4a and 4c. Here, HoPS ST gives the best results, because it always is in an agile mode, whereas ALE can only follow poor links quickly. Both RNP and F-LQE exhibit large variations in their ability to track link-quality. We are aware that F-LQE does not actually try to estimate the PSR, but from an application point of view, the PSR (or ETX) is the relevant link-quality metric. Because F-LQE has a very low correlation values, it cannot even track the relative behavior of the PSR, so that it is doubtful whether its values are useful for applications at all. A detailed view into the course of a few links reveals why LEEP is performing with a relatively low MAE: Many links have a quality-value close to 100% for long times. For these links, LEEP predicts a value of 0 ETX, yielding a 100% PSR and thus a low error. This partially cancels the heavy estimation errors when LEEP is not updating its quality metric. In contrast, pure EWMA approaches suffer from optimal links, since they are slow in taking extreme values.

Long-term link-quality estimation, see Fig. 4b and 4d, is handled well by the same candidates as before. ALE and all HoPS estimators except HoPS ST give very low MAE values with a median of about 1%. 75% of the links are predicted correctly with an error less than 2%. This is a lower value than the best quarter of LEEP estimates. It comes as a little surprise that LEEP shows a comparable performance for both short- and long-term ground truth traces. It therefore turns out, that for long-term link-quality assessment a strong low-pass behavior is beneficial, because many links had only short disturbances. However, this picture may slightly change for a different testbed.

The outliers in the CCF-plots are caused by very weak links with short bursts of medium or high PSR values. We did not filter these links out manually, as they are realistic phenomena and a clear criterion for removal could not be found. It shows that HoPS performs a little better in the presence of these links.

5.2 Link Selection

In many application scenarios, nodes have to select links from a set of available candidates. This could be the case for selecting the next hop in a routing protocol or for deciding whether to keep or replace a possible neighbor in a size-limited neighborhood table. The latter example is of particular importance in dense networks, in which nodes can only track a small subset of all available communication partners due to memory restrictions and the lack of dynamic memory allocation in most sensor node operating systems.

Figure 5 shows the link-selection decisions of the LQEs for all available links and after each 4 s time interval for thresholds of 85% and 95% link-quality, respectively. If the current quality estimate of a link is equal to or above the threshold, that link is selected (solid areas), otherwise it is rejected (hashed areas). The decisions of the LQEs are compared to the real course of the link for 1 min and 60 min windows. Correct decisions are the solid areas at the bottom (correct selections) and the hashed areas at the top (correct rejections) of each bar. False rejections are represented by the lower hashed areas and false selections

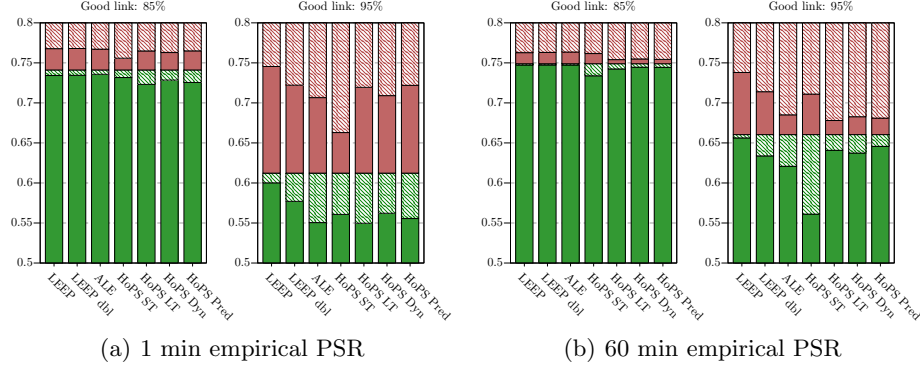


Fig. 5. Relative number of correct (solid) and wrong (hashed) link classifications (selections: bottom, rejections: top)

are indicated by the upper solid ones. The classification of the windows can also be read from the figures—e.g., Fig. 5(a) reveals that a fraction of 0.74 and 0.61, respectively, were above the two thresholds for the short-term window.

For the 1 min window in Fig. 5(a), HoPS ST achieves the smallest number of false decisions while HoPS LT generates the most errors due to its heavy smoothing. The other candidates produce comparable results for the 85% decision threshold. For the 95% one, LEEP makes a large number of false selections, because its resolution in this area is too low. Using a finer (LEEP dbl) resolution, this number is decreased at the cost of an increased amount of false rejections. In terms of cluster formation and long-term routing decisions, false selections may pose a severe hazard on throughput and network stability, whereas false rejections lead to a negative impact only in sparse networks with very few links. Here, network connectivity can be prevented, if there are too many false rejections. HoPS ST and Dyn achieve better selection results for the 95% decision threshold than ALE, where HoPS Dyn produces more false selections than HoPS ST.

The 60 min window in Fig. 5(b) flips the coin in favor of the long-term HoPS solutions HoPS LT and HoPS Pred. For the decision threshold of 85% they leave out some good links, but produce way less false selections than LEEP and ALE, which have the same performance in this setting. The 95% threshold gives an even more differentiated picture. HoPS Pred and HoPS LT make few wrong selections and rejections. HoPS Dyn produces comparable results; while falsely selecting almost as many links as ALE, the number of false rejections is only half as large. HoPS ST is too agile to produce good results. Increasing the resolution of LEEP decreases the number of false selections insufficiently.

To investigate the decisions of the LQEs in more detail, Fig. 6 portrays the relative distribution of selected (solid bars) and rejected (hashed bars) links for the long-term ground truth with 95% decision threshold. The solid curve indicates the Cumulative Distribution Function (CDF) for selected links, i.e., the fraction of selected links with an actual PSR on the X-coordinate. Similarly, the

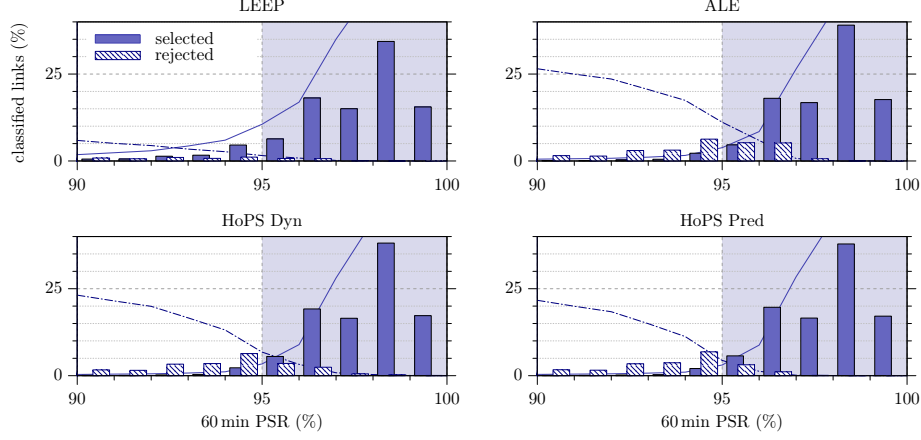


Fig. 6. Distribution of 60 min link PSRs for selected (solid) and rejected (hashed) links based on a decision threshold of 95% PSR (good links)

dash-dotted curve indicates the inverse CDF for rejected links, i.e., the fraction of rejected links with an actual PSR greater than the value on the X-coordinate. The crossing points at the 95% vertical borderline resemble the decision accuracy.

10% of the links selected by LEEP have an actual PSR less than 95%, while almost no links with a PSR of at least 97% have been rejected. The opposite case is true for ALE. More than 10% of the rejected links are actually better than 95%. At the same time, link selection works well: Less than 4% of the falsely selected links turn out to have a PSR below 94%. The selection CDF for HoPS Dyn and HoPS Pred follows almost the same course as the one of ALE, but the rejection CDF is lower—approximately 7% for HoPS Pred and 5% for HoPS Dyn. The two HoPS derivatives therefore exhibit a better rejection behavior of links below the decision threshold by incorporating the trend and variation. The main difference between HoPS Dyn and HoPS Pred is that the latter has identified the links in the region between 96% and 97% more accurately.

6 Conclusion

In this paper, we have presented a novel approach on link-quality estimation in wireless sensor networks: the Holistic Packet Statistics (HoPS). This new LQE calculates four distinct quality metrics, describing the short- and long-term quality of a link at the same time, while also providing information about the dynamics of a link by means of the variation and trend of the link-quality. All of these metrics can be accessed by the application. These metrics are calculated efficiently using EWMA filters and hence have a small memory footprint not exceeding that of well-known competitors, such as F-LQE or LEEP.

HoPS has been compared to a bouquet of existing LQEs by using PSR data gathered in an indoor testbed. Due to its four metrics, HoPS tracks the progress

of a link accurately, being therefore able to inform an application about the course of short- and long-term link-quality. An adaptive link-quality estimate can be provided by taking the variation and trend into consideration. This is supported by a comparative study of PSR traces of various links. An additional analysis of the correlation and average absolute error of real PSR values proves the gain of HoPS over existing LQEs. Moreover, it has been shown that HoPS improves link selection for routing decisions or neighborhood formation.

However, open issues and future work have been identified. Firstly, useful link-quality prediction is only possible, if link behavior is studied in more detail—e.g., sensing that a good link is currently deteriorating is more useful, if it is known in how many cases a link recovers quickly versus the times it breaks down completely. Secondly, the potential of HoPS can only be fully exploited, if its metrics are actually used by the application. Merging the four values into a single prediction value brings low improvements only. Therefore, advanced routing and neighborhood management protocols have to be devised. We plan to make further investigations on this ground, hoping to push link-quality assessment in wireless sensor networks a step forward.

7 Acknowledgments

Thanks to the anonymous reviewers for their insightful feedback. This material is based upon work partially supported by the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) under contract number TU 221/4-1.

References

1. Baccour, N., Koubaa, A., Youssef, H., Jamaa, M.B., do Rosario, D., Alves, M., Becker, L.: F-LQE: A Fuzzy Link Quality Estimator for Wireless Sensor Networks. In: EWSN (Feb 2010)
2. Becher, A., Landsiedel, O., Kunz, G., Wehrle, K.: Towards Short-Term Wireless Link Quality Estimation. In: Hot EmNets (Mar 2008)
3. Caleffi, M., Paura, L.: Bio-inspired Link Quality Estimation for Wireless Mesh Networks. In: WoWMoM (Jun 2009)
4. Cerpa, A., Wong, J.L., Potkonjak, M., Estrin, D.: Temporal Properties of Low Power Wireless Links: Modeling and Implications on Multi-Hop Routing. In: MobiHoc (May 2005)
5. Chen, Y., Terzis, A.: On the Mechanisms and Effects of Calibrating RSSI Measurements for 802.15.4 Radios. In: EWSN (Feb 2010)
6. Couto, D.S.J.D., Aguayo, D., Bicket, J., Morris, R.: A High-Throughput Path Metric for Multi-Hop Wireless Routing. In: MobiCom (Sep 2003)
7. Fonseca, R., Gnawali, O., Jamieson, K., Levis, P.: Four-Bit Wireless Link Estimation. In: HotNets VI (Nov 2007)
8. Gnawali, O.: The Link Estimation Exchange Protocol (LEEP) (Feb 2006), <http://www.tinyos.net/tinyos-2.x/doc/html/tep124.html>
9. Lal, D., Manjeshwar, A., Herrmann, F., Uysal-Biyikoglu, E., Keshavarzian, A.: Measurement and Characterization of Link Quality Metrics in Energy Constrained Wireless Sensor Networks. In: GlobeCom (Dec 2003)

10. Liu, T., Kamthe, A., Jiang, L., Cerpa, A.: Performance Evaluation of Link Quality Estimation Metrics for Static Multihop Wireless Sensor Networks. In: SECON (Jun 2009)
11. Senel, M., Chintalapudi, K., Lal, D., Keshavarzian, A., Coyle, E.J.: A Kalman Filter Based Link Quality Estimation Scheme for Wireless Sensor Networks. In: GlobeCom (Nov 2007)
12. Srinivasan, K., Levis, P.: RSSI is Under Appreciated. In: EmNets '06 (May 2006)
13. Weyer, C., Unterschütz, S., Turau, V.: Connectivity-aware Neighborhood Management Protocol in Wireless Sensor Networks. In: FGSN (Sep 2008)
14. Woo, A., Tong, T., Culler, D.: Taming the Underlying Challenges of Reliable Multihop Routing in Sensor Networks. In: SenSys (Nov 2003)