Efficient Online Estimation of Bursty Wireless Links

Muhammad Hamad Alizai, Hanno Wirtz, Georg Kunz, Benjamin Grap, Klaus Wehrle
Communication and Distributed Systems (ComSys), RWTH Aachen University, Germany
firstname.lastname@comsys.rwth-aachen.de

***Abstract*—Rapidly changing link conditions make it difficult to accurately estimate the quality of wireless links and predict the fate of future transmissions. In particular *bursty links* pose a major challenge to online link estimation due to strong fluctuations in their transmission success rates at short time scales. Therefore, the prevalent approach in routing algorithms is to employ a long term link estimator that selects only consistently stable links — PRR $> 90\%$ — for packet transmissions. The use of bursty links is thus disregarded although these links provide considerable additional resources for the routing process.**

Based on significant empirical evidence of over 100,000 transmissions over each link in widely used 802.15.4 and 802.11 testbeds, we propose two metrics, *Expected Future Transmissions* (EFT) and *MAC*$_3$, for runtime estimation of bursty wireless links. We introduce the Bursty Link Estimator (BLE) that, based on these two metrics, accurately estimates bursty links in the network rendering them available for packet transmissions.

I. INTRODUCTION

The availability and reliability of wireless links exhibits dynamic behavior at short and long time scales [1], [2]. Therefore, choosing the best link, in terms of routing progress and need for transmission resources, requires an accurate and timely estimation of the available links. Current link estimators, using metrics like Packet Reception Rate (PRR) and Expected Transmission Count (ETX), only capture link dynamics at long time scales for the sake of a stable routing topology. These metrics estimate the quality of a link over extended periods of time – in the order of minutes or hours – and thus achieve poor estimates for rapidly changing bursty links. As a result, bursty links are typically excluded from the routing process. However, recent protocol studies [3], [4] demonstrate that these links are long range and achieve significantly higher routing progress than stable links. Using these links therefore covers otherwise multiple transmissions and thus saves the energy and resource consumption coupled with these transmissions. Furthermore, studies show that typical traffic patterns in the Internet as well as in multihop wireless networks are bursty [5], [6]. Hence, an optimal online link estimation at the time of a burst benefits spontaneous transmissions as well as the overall network performance.

Link burstiness is a well established fact: it has been thoroughly analyzed [7], accurately modeled [2], and experimentally measured [8]. In this context, we define an intermediate wireless link with a PRR between 10% and 90% as a *bursty link*[1] if packet delivery on this link is correlated. This means that shifts between phases of reliable and poor packet delivery occur at short time scales, but future packet delivery is correlated to the recent success rate. Despite establishing a very strong knowledge base regarding the causes of link burstiness over the past few years, we still lack metrics that define the quality and usability of bursty links. Similarly, we need a link estimator that can assess link usability online (i.e., during runtime) to enable the inclusion of these links in the routing process. The definition of appropriate metrics and the design of a link estimator based on these metrics is our main contribution and departure from the existing work.

The requirements and challenges of estimating intermediate links are substantially different from conventional link estimation. For example, long-term packet reception rates (PRR) – otherwise the key link quality metric – of intermediate links do not suffice as a metric. Rather, (1) we are interested in whether or not packet delivery on an intermediate link is correlated to the recent delivery history, i.e., if the link is bursty or not. (2) We want to know how long a bursty link remains reliable for transmission, i.e., what is the length of successful transmission bursts. (3) We need to pinpoint exactly when a bursty link has a reliable or unreliable transmission period. None of these three pieces of information, which we consider keys to profitably using intermediate links for routing, are provided by existing link estimators.

This paper makes the following three contributions: First, we introduce MAC$_3$ as a metric to estimate the burstiness of links based on recent delivery traces. MAC$_3$ extends the established *Conditional Packet Delivery Function* (CPDF) [1] by calculating a moving average over the results of CPDF (*Moving Average CPDF*). Secondly, we define *Expected Future Transmissions* (EFT) as a metric to estimate the duration for which a bursty link remains reliable for transmission. We also show that EFT is strongly correlated to MAC$_3$. Finally, based on these two metrics, we introduce a Bursty Link Estimator (BLE), derive requisite parameters for its usage, and evaluate its efficacy in estimating intermediate links. Our results indicate that BLE identifies bursty links in the network with high accuracy, hence paving the way for such links to be included into the routing infrastructure.

[1]Stable links with a PRR $> 90\%$ can also be bursty but such links are not in the focus of this paper.

This paper is structured as follows. Firstly, we discuss related work Section in II. We then analyze and define the exact scope of our work in Section III. From this, we derive the design of our metrics and show their viability in Section IV. Finally, Section V presents the design and evaluation of our link estimator before we conclude in Section VI.

II. Related Work

Capturing link dynamics at different time scales and characterizing link burstiness have been the focus of many recent studies. We can divide prominent related efforts into three main categories.

Measuring Link Burstiness: In their seminal study on quantifying the extent and characteristics of bursty links, Srinivasan et. al. [8] define a factor β that measures the burstiness of a wireless link. β is calculated by using CPDFs [1], [8], which determine the success probability of the next transmission after n consecutive successes or failures. Hence, β is used to differentiate between bursty links with long bursts of successes or failures and links with statistically independent packet losses – with perfectly bursty and completely independent links marking the opposite ends of the spectrum. Although β is a very useful metric to measure link burstiness, its goal is the characterization of links based on existing traces rather than the online assessment of a link. Our evaluation in Section IV-C also reveals that calculating β over short history sizes, a fundamental requirement of online assessment, results in fluctuating and error-prone results.

Short Term Link Estimation (STLE): In our previous work on wireless link dynamics [3], [9], [10], we introduced a packet snooping based concept of STLE to analyze the impact of recent transmission success and failure rate on the future quality of a link at fine-grained time scales. However, although STLE is concerned with link estimation, we argue that the proposed mechanism only provides link discovery: STLE only tells whether or not a link becomes temporarily available but does not provide an estimation for how long this will be the case. Furthermore, no difference is made between recurring bursty links and accidental consecutive successful deliveries. These characteristics cause STLE to drop a link after just a single failed transmission, which further impedes its usability in real-world networks.

We developed Bursty Routing Extensions (BRE [3]) by significantly extending STLE. In BRE, any node that successfully overhears a limited number of consecutive transmissions from a neighboring node immediately declares that link reliable for transmission. Thus, multiple overhearing nodes may, at the same time, respond with an announcement that their link with the sender node is available. If these nodes then try to utilize the newly available link, an instant congestion of the wireless link ensues. As BRE is not able to predict whether the duration of an available link suffices for the intended traffic, there is no way of correctly estimating the usability of a link.

Long Term Link Estimation (LTLE): This is the traditional link estimation mechanism employed by the majority of current multihop wireless routing protocols [11], [12]. It is based on window mean exponential weighted moving averages (WMEWMA) of link PRRs or ETX [13]. Although this metric is highly accurate and has a small settling time for good and bad links, i.e., with PRRs close to 0% and 100%, it does not perform well for links of intermediate quality [14] – also indicated by our results in Section V-A. Hence, such link estimation mechanism cannot be used for estimating intermediate links at short time scales.

III. Problem Analysis

To provide a clear motivation for our work as well as a separation from the previously mentioned related work, we now define our problem space and the requirements for a solution. First, we introduce the prevalent network scenario to illustrate use cases and benefits of our solution. Second, we detail the motivation for employing a Bursty Link Estimator in networks where LTLE mechanisms are prevalent. Based on this, we highlight the key requirements of a link estimator for incorporating bursty links in the routing process.

A. Network Scenario

Wireless Sensor Networks (WSNs) and Wireless Mesh Networks (WMNs) provide flexible and robust ways of establishing network structures without the need for an exhaustive infrastructure. Routing structures in these networks are self-established and -maintained and depend on the presence of wireless links between nodes in the network. A resource-efficient utilization of these structures greatly increases throughput and network lifetime and reduces transmission energy and failures. Our work targets WSNs and WMNs due to their equivalent routing mechanisms.

B. The Need to Utilize Bursty Links

With regard to the characterization of links in [8], LTLE mechanisms typically utilize only good to perfect links with a PRR $\geq 90\%$. However, most links in wireless networks exhibit worse PRRs and are thus excluded from routing decisions. This results in i) a stable and clear-cut routing topology, ii) usage of short range links and little routing progress on each hop and iii) heavy utilization of the selected links. In LTLE, a trade-off is thus made between the high cumulative resource consumption of series of short range links and the ease of utilizing only a fraction of the existing links. In contrast to this, bursty long-range links offer high routing progress with only one transmission but need to be included in the routing process. This inclusion requires an accurate online estimation of bursty links which is not possible using prevalent link estimators.

In wireless multihop networks, such as WMNs and WSNs, the networking hardware is the most dominant consumer of energy. The amount of energy consumed by the networking hardware is directly proportional to the number of transmissions required by a packet to reach its destination. By utilizing bursty links with significantly better routing progress (i.e. less number of hops traversed) [3], [4], the number of transmissions and thus the amount of energy consumed is reduced.

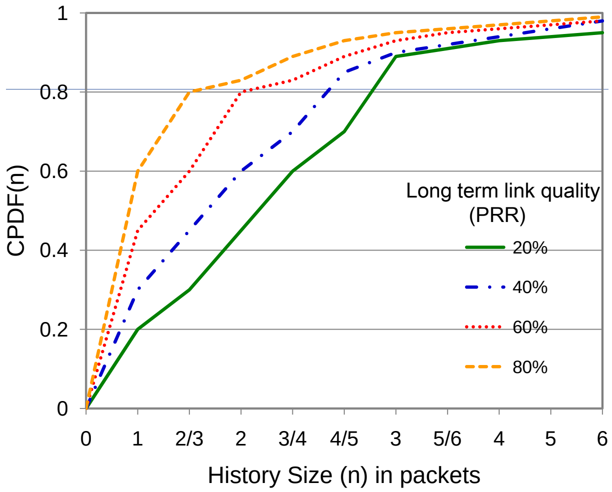

Fig. 1. Measuring the impact of recent transmission success over a link on the next transmission over that link. A label of k/n stands for k successes during the last n transmissions, and n is a shorthand for n/n. CPDF(n) is the probability that the next transmission is successful..

C. Requirements of a Bursty Link Estimator

The design of a link estimator that reliably reflects the state of a given link has to fulfill multiple requirements. First, appropriate metrics need to be derived as key building blocks of the estimator. Such metrics must timely estimate the current link quality based on a very short transmission history, in order to adapt to the rapidly changing conditions of bursty links. Additionally, the predicted link quality must accurately and reliably reflect the actual link quality, i.e., the estimation error needs to be small and stable.

Second, building upon such metrics, a link estimator must efficiently utilize the given information to select beneficial links for routing. This requires appropriate neighbor table management policies that select those links for routing – among all available links – which allow for the best routing progress.

IV. Deriving Metrics for Bursty Links

Based on the properties specified in the previous section, this section defines and evaluates two metrics, MAC_3 and EFT, that i) identify bursty links in the network and ii) estimate the length of successful transmission bursts. These metrics subsequently lay the foundation for BLE. In the following, we first provide detailed information on the particular data set used in the remainder of this study.

A. Data Set and Experimental Model

The design of our link quality metrics as well as the resulting BLE are based on widely used empirical data rather than a theoretical model. We strongly believe that empirical observations from multiple real world scenarios are important both for developing metrics and evaluating the efficacy of the concepts presented in this paper. This is because the transmission fluctuations and dynamics revealed by bursty links of intermediate quality are hard to capture in a theoretical model only [1], [2], [8] .

The evaluation results presented in this section are therefore based on the SING mesh data-set [15] compiled at Stanford University and used in many recent state-of-the-art studies [8], [16], [17] on wireless link dynamics. It is a comprehensive data set collected from multiple IEEE 802.11 and IEEE 802.15.4 testbeds[2] including both packet and byte level radios such as the cc2420 and cc1000. Unless otherwise noted, our study utilizes the data provided by the Mirage testbed [19] – a 100 node micaz sensornet testbed at Intel Research Berkeley. The nodes are spread out over an indoor area of approximately 160' by 40'. Specifically, the data comprises traces of transmissions on channel 26 at a transmission power level of 15 dBm. Each node broadcast a burst of 100,000 packets with a packet inter arrival time of 10 ms. Out of all available links, we only include intermediate links in our evaluation and comparison. This is because good links would measure highly in our metrics and would thus improve our results. However, these links are not the focus of our work.

B. Case Study: Predicting Transmission Success from a Short History

Before introducing our metrics MAC_3 and EFT we motivate our goals and approach with a case study. We address the question, whether a short history of successful transmissions is sufficient to predict with a high probability that the next transmission on this link will be successful, too. Figure 1 depicts the conditional probability of a successful packet transmission based on the average long-term link quality (i.e., PRR) and a short-term history of consecutively *successfully* transmitted packets. It shows that for a link with a long term quality greater than 60%, even a single or two successful transmissions over that link raise the success probability of the next transmission to 90%. Similarly, it shows that for any link, regardless of its long term link quality, the probability of a future successful transmission is greater than 90% if the last three packets over that link were sent successfully.

Overall, these results validate our previous observations [3], and indicate that a short-term history of three packets over a link is sufficient to determine with a high probability whether the next transmission will be successful or not. In the following, we introduce two metrics MAC_3 and EFT that determine the success probability of future transmissions on a per link granularity, allowing us to reflect spatial properties of link dynamics.

C. Online Estimation of Link Burstiness

Estimating the burstiness of a link is mandatory to determine whether or not an intermediate link is beneficial to the overall routing performance. The key challenge is to clearly distinguish intermediate links with correlated packet losses from those with independent losses. However, unlike offline measurement mechanisms like β, we are not interested in how close a link is to an ideal bursty link with one long burst of

[2]Motelab [18], Mirage [19], and SWAN testbeds. Please visit http://sing.stanford.edu/srikank/datasets.html and the websites of each testbed for further information, e.g. topology, connectivity.

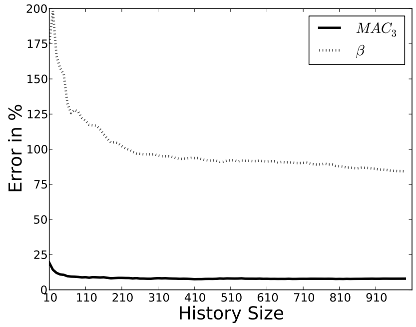

(a) Settling Time: MAC_3 shows a faster convergence towards its base value over the history size and achieves a smaller estimation error (7%) than β. Based on data of the Mirage testbed.

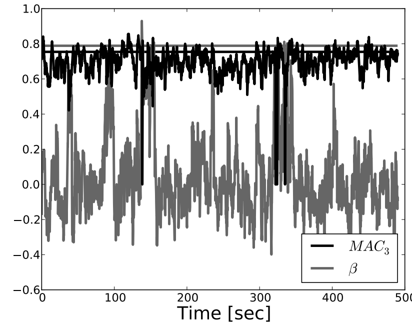

(b) MAC_3 generates more accurate and more stable results over time than β. The straight lines show the base values of both metrics over the entire transmission trace.

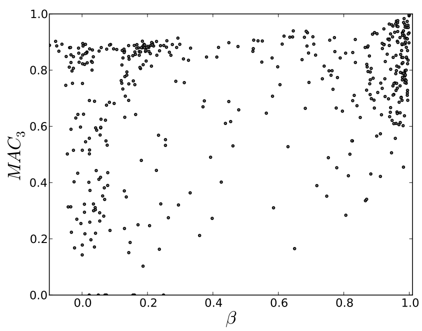

(c) MAC_3 reveals that many links have a high probability for a further successful transmission after three consecutive deliveries even though their β is very low.

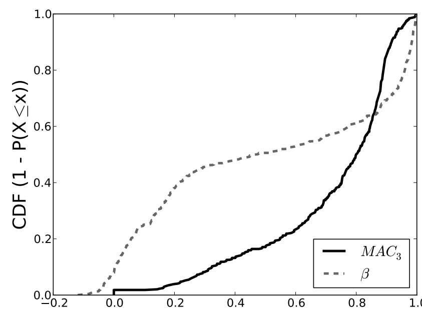

(d) Cumulative distribution of intermediate links: The majority of intermediate links is bursty ($\text{MAC}_3 > 0.7$), offering useful transmission opportunities.

Fig. 2. Comparing MAC_3 and β as a link burstiness metric for runtime link estimation. We use a smaller version of β for online link assessment. Our version of β does not enforce a confidence interval of 95% for its data points.

either successes or failures. Our goal to predict link burstiness at runtime strongly influences the definition of burstiness and the timescale of our prediction. In this context, we define a link as bursty as long as we can *recurrently* predict the fate of only the next transmission over a link with high probability. This is why we introduce a new metric that monitors a link for a limited transmission history and expresses if the occurrence of a successful transmission burst over a particular link is a mere coincidence or if it is a reoccurring trend. This information is important to determine if a link is beneficial for routing purposes.

Our online metric Moving Average $CPDF$ (MAC_3) is based on a $CPDF(n)$ (Conditional Packet Delivery Function) [1] which calculates the probability of one successful transmission following n previously successful transmissions. Based on the results in the initial case study in Section IV-B, we compute an average $CPDF(3)$ over the recent history h of length m of a link and denote it AC_3 (see Equation 1). $|CPDF(3)|$ defines the number of valid $CPDF(3)$ in the history. We define MAC_3 as the moving average of AC_3 that is computed by adding new values and removing old ones from the history at runtime.

$$AC_3 = \frac{\sum_{i=1}^{m} CPDF_{h_i,h_{i+1},h_{i+2}}(3)}{|CPDF(3)|} \tag{1}$$

To evaluate MAC_3 we compare it with the β factor [8] because, (i) it is the only metric available that measures link burstiness and, (ii) it enables a better understanding of the effectiveness of MAC_3 as a runtime metric. However, this comparison, by any means, does not attempt to undermine the usefulness of β as it was never developed for runtime measurements.

Figure 2(a) illustrates the estimation error of MAC_3 and β over history sizes ranging up to 1000 packets. The estimation error is the difference between the estimated value of either metric when applied to a certain history size (plotted on the x-axis) and the value when applied to the whole transmission trace (i.e., base value). The base value of β is calculated according to the procedure prescribed in [8]: a $CPDF(n)$ for a certain n is only considered in β calculations if it has at least 100 data points to achieve a 95% confidence interval of [p-0.1, p+0.1]. Whereas, for calculating β over a shorter transmission history, we do not enforce the condition of 100 data points. This is because, (i) it is simply not possible to collect 100 data points in a shorter transmission trace, and (ii)

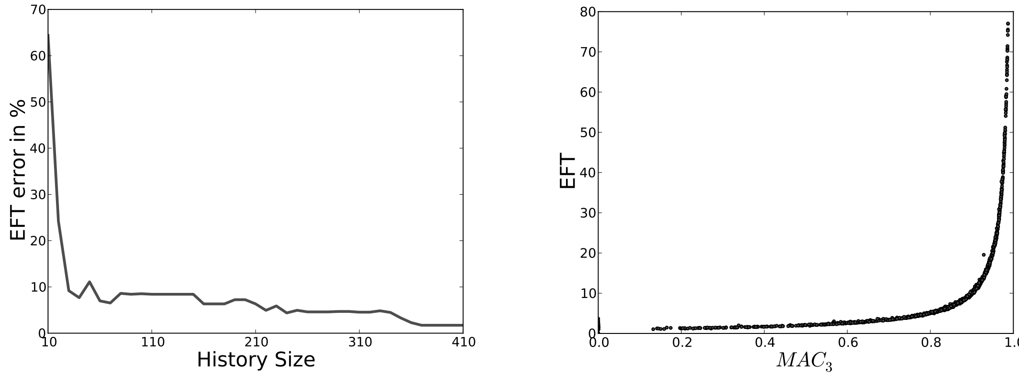

(a) Influence of the history size on convergence of EFT.

(b) Correlation between MAC_3 and EFT.

Fig. 3. EFT and MAC_3 as link quality metrics

we want to investigate if this restricted version of β provides accurate estimates and can be used for runtime estimation of link burstiness. The figure indicates that our online metric MAC_3 rapidly converges to a minimal error of 7% with a history size of less than 100 packets. In contrast, β shows a significantly slower initial convergence phase and is not able to achieve an error smaller than 83% even with a history size of 1000 packets. Moreover, β is not able to provide stable results for small history sizes as shown in Figure 2(b). Given a concrete history size, β generates severe fluctuations in its output over time when applied to an entire transmission trace of a particular link. The estimated values of MAC_3 on the other hand expose considerably smaller differences. In addition, the results of β again strongly deviate from the base value calculated over the whole trace (straight gray line) while the estimates of MAC_3 oscillate around its actual base value (straight black line). Overall, these results show the efficiency of MAC_3 as online metric: it is stable for short history sizes.

Next we show that MAC_3, in contrast to β, captures the short term properties of a link. Figure 2(c) shows that many links with $\text{MAC}_3 > 80\%$ have low β values. It means that on such links the probability of a successful transmission after three consecutive deliveries is greater than 80%, but the use of β as a link metric will not let a routing protocol select this link. After evaluating the effectiveness of MAC_3, we need to analyze what proportion of the available intermediate links are actually useful for routing. Figure 2(d) shows the cumulative distribution function of MAC_3 and β for all intermediate links in the Mirage testbed. We can clearly observe that the majority of these links have a very high MAC_3. As a result, MAC_3 unlocks the formerly wasted potential of those links and enriches the routing process with a multitude of new routing opportunities.

Concluding, MAC_3 is a lightweight metric for estimating link burstiness during runtime. Our results in Section V-C demonstrate that, when used as metric to estimate link burstiness, MAC_3 accurately identifies bursty links in the network.

D. *Estimating Burst Lengths*

In addition to identifying whether or not a link is bursty, a second metric for estimating the length of bursts is required. To illustrate why, assume a bursty link with a steady rate of bursts covering four successful transmissions each before becoming unreliable again. Such a link exhibits a very high $CPDF(3)$ value, causing MAC_3 to correctly identify it as bursty. However, if selected for transmission, this links allows for only one more successful transmission per burst, hence rendering it barely suitable for routing.

To solve this problem, we introduce a new metric named $EFT(n)$: It estimates the number of r successful future transmissions after n successful packet deliveries. This metric thus predicts the length of bursts and allows the link estimator to identify bursts of relevant size. Just like MAC_3, EFT uses an averaging moving window to traverse a transmission history. For each occurrence of three successful consecutive transmissions in the history, EFT determines the number of subsequently following successful transmissions and incorporates it in the total average.

Similar to MAC_3, EFT has a very small settling time (cf. Figure 3(a)): It converges to within 10% error at a history size of approximately 100 packets. Figure 3(b) indicates a strong correlation between EFT and MAC_3. For values of MAC_3 in the range of 0.1 to 0.7, EFT predicts burst lengths not longer than five packets. However, when MAC_3 exceeds 0.7, the estimated burst lengths increase significantly. As a result, we derive a threshold of 0.7 of MAC_3 for a link to be considered useful for future transmissions.

V. THE BURSTY LINK ESTIMATOR

We propose a packet snooping based link estimation [11], [14] for BLE. BLE is not supposed to work independently: it is an additional component of the routing infrastructure that enables fine-grained estimation of intermediate links and allows for such links to be included in the routing process. In this section, we first discuss why PRR is not a suitable metric for intermediate links and propose a combination of MAC_3 and EFT to be used as link quality metrics for BLE. We then

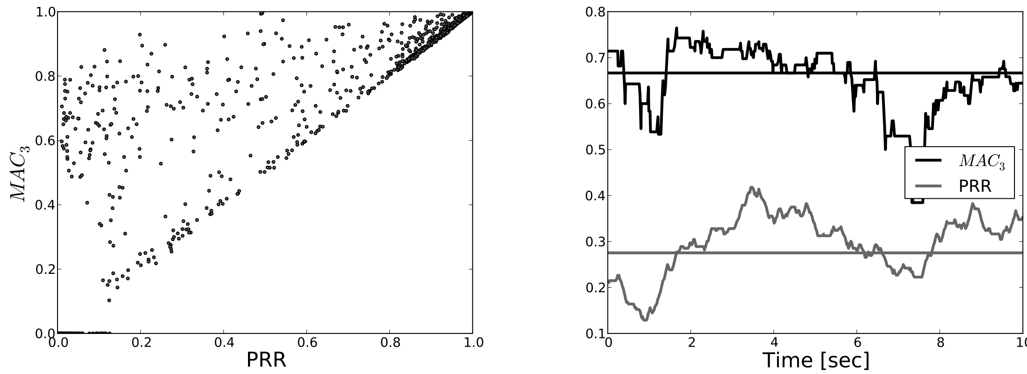

(a) Many links with low PRR exhibit a high success probability after 3 successful deliveries. All the links in Mirage testbed are plotted in this graph.

(b) A 28% link shows a high MAC_3 and stable progress over time. PRR will never include such a link in neighbor tables.

Fig. 4. Comaprison of MAC_3 with PRR.

provide further details about the information maintained in BLE's table. Finally, we conclude this section by evaluating BLE.

A. Link Quality Metric

PRR (or ETX: the reciprocal of PRR) is commonly used as a link metric in current link estimators. The basic technique is to calculate weighted moving averages of PRR over a very long time period. Similar to β, PRR does not fulfill the desired properties of a metric for our envisioned link estimator. For example, it is unable to capture short term dynamics exposed by bursty links of intermediate quality. Figure 4(a) highlights this fact: many links with a very high MAC_3 have very low PRRs. This means that over a long time scale these links have bad reception rates. However, when observing a history size as small as 3, it is possible to predict the success of future transmissions with high probability. Similarly, Figure 4(b) supports this argument by comparing PRR and MAC_3 over time. It shows that although MAC_3 indicates a high probability of successful delivery, PRR is unable to capture this reliable transmission period of a link. Hence, the use of PRR prohibits the exploitation of bursty links that offer useful transmission opportunities at shorter time scales.

In BLE, we use a hybrid metric that is based on the product of MAC_3 and EFT. Both MAC_3 and EFT are calculated by applying a sliding window over the packet delivery history of size h for each link in the table. Since maintaining the link history is an expensive memory operation and impacts the scalability, it is important to choose the threshold h appropriately as discussed in Section V-C1.

B. Table Management

BLE follows the basic table management algorithm outlined by Woo et. al. [14] and used by the majority of current link estimators [11], [20]. We deviate from the established concept in terms of (1) link selection as BLE only estimates intermediate links (i.e. PRR $<$ 90%), specifically, the ones not present in the table of the associated long term link estimator, and (2) different ingredients for the link insertion, eviction, and reinforcement policies. The estimator maintains a small table (e.g., of size 10) of candidate links which holds the following information per link:

MAC_{3in}: The reception MAC_3 of the link (1 byte).
EFT_{in}: The reception EFT of the link (1 byte).
MAC_{3out}: The sending MAC_3 of the link (1 byte).
EFT_{out}: The sending EFT of the link (1 byte).
Link History: The packet delivery history of size h (16 bytes, cf. Section V-C1). Bit arrays are used with 1 representing a successful delivery and 0 representing a failed transmission.
Available: A flag to determine if the link, with MAC_3 and EFT above certain threshold, is currently available for transmission. Set to 1 if the last three transmissions were successful over the link, and 0 otherwise.
Valid: A flag to determine if the link has a large enough delivery history, and all other table entries are up-to-date.

The table management is concerned with three tasks: Adding links, deleting links and maintaining links in the table. A new link is added to the table upon reception of a packet on a non-resident link and (i) a vacant entry in the table exists, (ii) the product of MAC_3 and EFT of a resident link drops below a user-specified threshold, or (iii) an entry expired due to a broken link or an insufficient packet reception rate. Additionally, link maintenance is performed after i received packets. At this point, all entries in the table are recalculated. The value i is a trade-off between the computational overhead and actuality of BLE.

C. Evaluation

We have implemented a prototype of BLE in TinyOS for a notoriously difficult class of wireless mesh networks – sensornets. Our evaluation of BLE focuses on two factors: (i) *Link History Size*: We empirically derive a requisite history size h that shall be maintained by BLE to compute its link metrics. (ii) *Link Estimation*: We validate that BLE indeed includes bursty links of high quality in the neighbor table. The

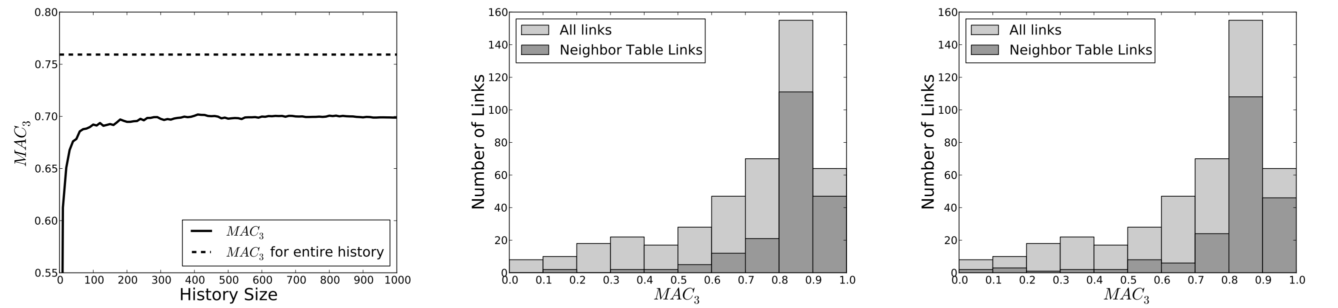

(a) Influence of the history size on convergence of MAC_3.

(b) Number of links with a given estimated quality (light gray); subset of these links that are included in the neighbor table after 1000 transmissions by each node in the network (dark gray).

(c) Number of links with a given estimated quality (light gray); subset of these links that are included in the neighbor table after 2000 transmissions by each node in the network (dark gray).

Fig. 5. Evaluating BLE

latter constitutes the key factor in assessing the performance of any link estimation mechanism, as the quality of the link selection process has a significant impact on the overall routing efficiency. For our evaluation we only use intermediate links from the network.

1) Link History Size: Although determining link history is a user-desired accuracy threshold, we derive their values here for completeness and for evaluation purposes. Our goal is to find a requisite history size that balances estimation error and memory consumption. A too small history does not provide enough information to enable BLE to accurately predict the link quality. Conversely, a too large history blocks valuable system resources and potentially does not even improve prediction accuracy. We assume that an estimation error of 10% yields user-acceptable results. Figure 2(a), 3(a) and 5(a) show our results derived from the Mirage testbed data. We clearly observe that MAC_3 and EFT converge below a 10% error at a history size of approximately 100 packets. Hence, for our evaluation, we use 16 bytes of memory to store a single link history which corresponds to a (large-enough) history size of 128 packets.

2) Link Estimation: This evaluation aims to confirm that BLE correctly identifies bursty links in the network to provide these links with a high value of MAC_3 for inclusion in the routing process. Figures 5(b) illustrates the total number of links with a certain estimated quality and the fraction of links that were included in the neighbor table by BLE after 1000 transmissions over each link in the network. We observe that the fraction of selected links increases in conjunction with the estimated link quality. The fact that not all links with a high value of MAC_3 are included for routing stems from the criteria of Link Addition (see Section V-B) and the requirements of a fixed and small table size so that there may exist more suitable links than can be included in the table. Although Figures 5(b) presents an instantaneous snapshot of the BLE tables in the network, we observed a similar trend throughout our evaluation.

3) Routing: Although an advanced routing evaluation is a future work and not part of our main research contribution in this paper, we present initial results here for completeness. Hence, our goal is not to design an optimized routing protocol but merely to integrate BLE with an existing routing protocol and link estimator to assess its potential benefits. We integrated BLE with the standard Collection Tree Protocol (CTP) [12] and the 4 Bit Link Estimator [20] shipped with TinyOS. By integrating BLE with CTP, we allow CTP to use long range intermediate links whenever (i) BLE declares a bursty link reliable for transmission, (ii) the MAC_3 of that link exceeds the predefined threshold, and (iii) the declared link offers a shorter routing path than the link currently used by CTP (i.e., by comparing their hop counts to the collection root). We randomly selected 5 node pairs[3] from MoteLab [18] as senders and collection roots. The maximum path length between these node pairs is 5 hops.

We want to analyze three factors: (a) How many intermediate links (the links disregarded by CTP) are taken by a packet on its path from source to destination (see x-axis in Figure 6), (b) What is the length of successful transmission bursts over these intermediate links used by CTP (see y-axis), and (c) How often these successful transmission bursts occur on an intermediate link (see the width of the circles). This third factor is important to observe if an intermediate link becomes repeatedly reliable for transmission or if a successful transmission burst over this link is a mere coincidence. Figure 6 shows that BLE enables routing protocols to use the previously ignored class of intermediate links with longer successful transmission bursts. It also shows that a packet takes multiple intermediate links on its way from source to destination. Moreover, we can clearly see that these links become repeatedly reliable for transmission as indicated by the radius of circles. Hence, these results prove the principle feasibility of BLE for routing.

VI. Conclusions and Future Work

We presented a bursty link estimator that allows the inclusion of bursty links into the routing process, thereby enabling

[3]Please visit http://motelab.eecs.harvard.edu/ to see the exact location of the selected node pairs and the overall network topology

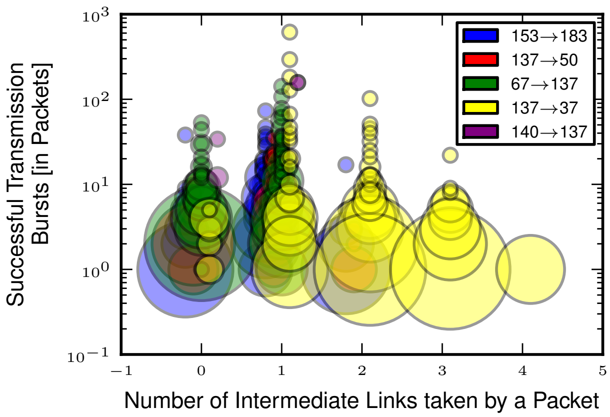

Fig. 6. The number of bursty links taken by a packet and the burst length on the path from source to destination. A randomly selected set of node-pairs (see legend) is used from MoteLab as senders and collection roots. The radius of the circle shows the number of occurrences of such transmission bursts. Please note the logarithmic y-axis.

a better utilization of the existing links in a network. We observed that the traditional metrics, β and PRR, used to measure link burstiness and link quality are of limited use in estimating intermediate wireless links. In this regard, we presented MAC_3 and EFT as metrics to estimate link burstiness and burst lengths of intermediate links, respectively. Our evaluation validates that a link estimator based on these two metrics accurately estimates intermediate links and enables inclusion of bursty links in the neighbor table.

Integrating BLE with routing protocols in different wireless domains — IEEE 802.11 and IEEE 802.15.4 — and a detailed evaluation of its performance benefits is ongoing work. Similarly, we are interested in evaluating the generality of our parameters, such as history size and error thresholds, and using more rigorous approaches for selecting these parameters. Moreover, our interest lies in understanding how BLE workscales with different table sizes, node densities, topologies and traffic patterns.

Acknowledgments

Many thanks to Kannan Srinivasan and Philip Levis for providing access to their large collection of data sets comprising measurements from both 802.15.4 and 802.11 testbeds. This research was funded in part by the DFG Cluster of Excellence on Ultra High-Speed Mobile Information and Communication (UMIC), German Research Foundation grant DFG EXC 89.

References

[1] H. Lee, A. Cerpa, and P. Levis, "Improving wireless simulation through noise modeling," in *IPSN '07: Proceedings of the 6th international conference on Information processing in sensor networks*. ACM, 2007, pp. 21–30.

[2] A. Kamthe, M. A. Carreira-Perpi nán, and A. E. Cerpa, "M&m: multi-level markov model for wireless link simulations," in *SenSys '09: Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*. ACM, 2009, pp. 57–70.

[3] M. H. Alizai, O. Landsiedel, J. A. B. Link, S. Götz, and K. Wehrle, "Bursty traffic over bursty links," in *SenSys '09: Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*. ACM, 2009, pp. 71–84.

[4] S. Biswas and R. Morris, "Exor: opportunistic multi-hop routing for wireless networks," in *SIGCOMM '05: Proceedings of the 2005 conference on Applications, technologies, architectures, and protocols for computer communications*. ACM, 2005, pp. 133–144.

[5] J. Zander, "Radio resource management in future wireless networks: requirements and limitations," *Communications Magazine, IEEE*, vol. 35, no. 8, pp. 30 –36, aug. 1997.

[6] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis, "Some implications of low power wireless to ip networking," in *Proceedings of the 5th Workshop on Hot Topics in Networks (HotOS V)*, 2006.

[7] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*. ACM, 2004, pp. 121–132.

[8] K. Srinivasan, M. A. Kazandjieva, S. Agarwal, and P. Levis, "The β-factor: measuring wireless link burstiness," in *SenSys '08: Proceedings of the 6th ACM conference on Embedded network sensor systems*. ACM, 2008, pp. 29–42.

[9] M. H. Alizai, O. Landsiedel, K. Wehrle, and A. Becher, "Challenges in short-term wireless link quality estimation," FU Berlin, Germany, Tech. Rep., 2008.

[10] A. Becher, O. Landsiedel, G. Kunz, and K. Wehrle, "Towards short-term link quality estimation," in *Proceeding of the 5th workshop on Embedded Networked Sensors (HotEmNets)*, 2008.

[11] R. Fonseca, S. Ratnasamy, J. Zhao, C. T. Ee, D. Culler, S. Shenker, and I. Stoica, "Beacon vector routing: scalable point-to-point routing in wireless sensornets," in *NSDI'05: Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation*. USENIX Association, 2005, pp. 329–342.

[12] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis, "Collection tree protocol," in *SenSys '09: Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*. ACM, 2009, pp. 1–14.

[13] D. S. J. De Couto, D. Aguayo, J. Bicket, and R. Morris, "A high-throughput path metric for multi-hop wireless routing," *Wirel. Netw.*, vol. 11, no. 4, pp. 419–434, 2005.

[14] A. Woo, T. Tong, and D. Culler, "Taming the underlying challenges of reliable multihop routing in sensor networks," in *SenSys '03: Proceedings of the 1st international conference on Embedded networked sensor systems*. ACM, 2003, pp. 14–27.

[15] SING Datasets, Standford University, 2009, Stanford Information Networking Group. [Online]. Available: http://sing.stanford.edu/srikank/datasets.html

[16] K. Srinivasan, M. Jain, J. I. Choi, T. Azim, E. S. Kim, P. Levis, and B. Krishnamachari, "The κ-factor: Inferring protocol performance using inter-link reception correlation," in *Proceedings of the 16th Annual International Conference on Mobile Computing and Networking (Mobicom 2010)*. ACM, 2010.

[17] M. Wachs, J. I. Choi, J. W. Lee, K. Srinivasan, Z. Chen, M. Jain, and P. Levis, "Visibility: a new metric for protocol design," in *SenSys '07: Proceedings of the 5th international conference on Embedded networked sensor systems*. ACM, 2007, pp. 73–86.

[18] G. Werner-Allen, P. Swieskowski, and M. Welsh, "Motelab: a wireless sensor network testbed," in *Proceedings of the 4th international symposium on Information processing in sensor networks*, ser. IPSN '05. Piscataway, NJ, USA: IEEE Press, 2005. [Online]. Available: http://portal.acm.org/citation.cfm?id=1147685.1147769

[19] B. N. Chun, P. Buonadonna, A. Auyoung, C. Ng, D. C. Parkes, J. Shneidman, A. C. Snoeren, and A. Vahdat, "Mirage: a microeconomic resource allocation system for sensornet testbeds," in *Proceedings of the 2nd IEEE Workshop on Embedded Networked Sensors (EmNets)*, May 2005.

[20] R. Fonseca, O. Gnawali, K. Jamieson, and P. Levis, "Four-bit wireless link estimation," in *Proceedings of the Sixth Workshop on Hot Topics in Networks (HotNets-VI)*, 2007.