

Bursty Traffic over Bursty Links

Muhammad Hamad Alizai, Olaf Landsiedel, Jó Ágila Bitsch Link, Stefan Götz, Klaus Wehrle
Distributed Systems Group, RWTH Aachen University, Germany

{hamad.alizai,olaf.landsiedel,jo.bitsch,stefan.goezt,klaus.wehrle}@rwth-aachen.de

Abstract

Accurate estimation of link quality is the key to enable efficient routing in wireless sensor networks. Current link estimators focus mainly on identifying long-term stable links for routing. They leave out a potentially large set of intermediate links offering significant routing progress. Fine-grained analysis of link qualities reveals that such intermediate links are *bursty*, i.e., stable in the short term.

In this paper, we use short-term estimation of wireless links to accurately identify short-term stable periods of transmission on bursty links. Our approach allows a routing protocol to forward packets over bursty links if they offer better routing progress than long-term stable links. We integrate a Short Term Link Estimator and its associated routing strategy with a standard routing protocol for sensor networks. Our evaluation reveals an average of 19% and a maximum of 42% reduction in the overall transmissions when routing over long-range bursty links. Our approach is not tied to any specific routing protocol and integrates seamlessly with existing routing protocols and link estimators.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless communication; C.2.2 [Network Protocols]: Routing protocols

General Terms

Algorithm, Design, Experimentation, Performance

Keywords

Bursty Links, Routing, Link Estimation

1 Introduction

Instability of links and connectivity in low-power wireless sensor networks (WSNs) has so far been regarded as a difficult problem that existing routing algorithms try their utmost to avoid. Therefore, since the emergence of WSNs,

research has mainly focused on link estimation and routing techniques [7, 8, 22, 31] which identify and utilize consistently high quality links for packet forwarding. Links of intermediate quality¹ are ignored to ensure routing stability and to attain high end-to-end reliability. Protocol studies [1, 14, 24, 26] have shown that these intermediate quality links are bursty, i.e., they frequently switch between stable and unstable periods of transmission for a limited number of consecutive packets. In this paper, we argue that: (1) Bursty links can be used for packet forwarding during their stable periods without affecting the reliability and stability of existing routing protocols; (2) These links often achieve significantly better routing progress and routing throughput than the long-term links chosen by existing routing protocols.

Today's link estimators [7, 8, 31] measure the quality of a link in the ETX metric: the number of (re)transmissions required for a successful transmission. Typically, link estimators periodically broadcast control packets, so-called beacons, to maintain up-to-date ETX estimates. Widespread routing protocols in WSNs, such as BVR [8] and CTP [10, 22], select links as suggested by their link estimator. To achieve better connectivity and reliable packet communication, today's link estimators restrict communication to neighbors with constantly high-quality links. These links are identified based on the long-term success rate of a link collected over a time frame in the order of minutes. However, this approach has two major pitfalls. First, neighbors with intermittent connectivity might reach farther into the network. Their use might therefore offer better routing progress and hence reduce the number of transmissions, lower energy usage in the network, and increase throughput. Second, in a sparse network with low density of nodes, a node might have no high-quality neighbor in its communication range, requiring a mechanism to deal with unstable connectivity.

1.1 Significance and Distinction

A great deal of effort has been invested in evaluating and quantifying the varying and dynamic characteristics of links in WSNs both analytically [14, 16] and experimentally [1, 6, 27]. To fine-tune protocol parameters operating at different layers of the network stack, these studies have led to the definition of analytical metrics and experimental parameters, such as link burstiness [26]. Our major depart-

¹we use the term intermediate quality to represent wireless links with a PRR between 10% and 90% [26].

ture from the existing work is that we neither introduce any new experimental model for wireless links nor define any parameters for fine-tuning protocols. We exploit the existing knowledge about burstiness of wireless links to enhance routing performance. Similarly, we investigate the usability, applicability, and practicality of routing over bursty links that often offer the highest routing progress [5,31], instead of devising mechanisms to bypass them for the sake of instability.

In order to utilize these intermediate-quality links in the routing process, we use short-term link estimation that captures link dynamics at a high resolution in time. It identifies the periods when bursty links become temporarily reliable or unreliable for transmission. However, to maintain a stable network topology, we do not replace existing link estimators. We show the seamless integration of a Short-Term Link Estimator (STLE) and an adaptive routing strategy with existing routing protocols and link estimators.

1.2 Contribution

Overall, this paper has three key contributions. First, it shows how short-term link estimation can be used for fine-grained estimation of bursty links to identify stable transmission periods. Thereby it enables routing protocols to forward packets over long-range bursty links and minimize the number of transmissions in the network. Second, we present an adaptive routing strategy which uses STLE for packet forwarding over bursty links. Third, we present Bursty Routing Extensions (BRE) - an integration of STLE and the adaptive routing strategy with a standard routing protocol for sensor networks. As a result, we show how our approach can be integrated with existing routing protocols and link estimators.

We evaluate the performance of BRE on two different testbeds: (1) MoteLab [2] at Harvard University and (2) TWIST [12] at Technical University of Berlin. Our evaluation measures that BRE achieves an average of 19% and a maximum of 42% reduction in the number of transmissions when compared to a traditional collection protocol i.e. CTP [10].

The remainder of this paper is structured as follows. Section 2 discusses related work. Section 3 provides an overview of the basic concepts of STLE and derives design goals. The design of BRE and the associated challenges are discussed in detail in Sections 4. We present our evaluation results in Section 5. Section 6 concludes our discussion and outlines future work.

2 Related Work

In their seminal study, Woo et al. [31] investigated reliable multi-hop routing in wireless sensor networks. Their recommendations include the use of PRR as link quality metric for link estimators, interpreting the minimum data transmission rate as a link availability indicator, and using the number of transmissions from source to destination as a routing metric. Among the well established estimation techniques [30] examined in their study, a window mean estimator with exponentially weighted moving average (WMEWMA) performed the best overall. Although WMEWMA is highly accurate and has a small settling time for good and bad links, i.e., close to 0% and 100% PRR, it does not perform well for links of intermediate quality [3,31]. The four-bit link estima-

tor [7] extends WMEWMA by combining information from the network-, link-, and physical layers to reduce packet delivery cost.

2.1 Packet Loss is Correlated

The majority of existing link estimation techniques assume that individual packet loss events on a link are statistically independent of each other and that they follow a Bernoulli distribution [30]. However, studies such as [3, 6, 25] mark this assumption as inappropriate when wireless links are estimated over shorter time scales. For example, in our previous work [3], we analyze the impact of recent transmission success and failure rate on the future quality of a link at fine-grained time scales. We associate the recent history of successful transmission and failure to predict the short-term quality of a link, i.e., its availability for transmission. The conclusion of this study is that any link, no matter of what quality, becomes temporarily reliable after h consecutive packets are received over that link. The value of h is determined to be about three packets for the links of intermediate and four to five packets for the links with bad quality. Another finding is that the fraction of intermediate links increases with distance, i.e., the probability to find a long distance link of intermediate quality is higher than the probability to find a good link.

2.2 Link Burstiness

Srinivasan et al. [26] present a comprehensive study to quantify the extent and characteristics of bursty links. They define a factor β , which measures the burstiness of a wireless link. β is calculated by using conditional probability distribution functions (CDFs), which determine the probability that the next packet will be received after n consecutive successes or failures. β is used to identify bursty links with long bursts of successes or failures and statistically independent links, with ideal bursty ($\beta = 1$) and independent ($\beta = 0$) links marking the two ends of spectrum. To explore the feasibility of β , an *opportunistic* transmission algorithm is introduced. The underlying idea is to pause transmissions after experiencing packet failure over a good quality link, with packet losses in bursts, to reduce the number of transmissions in the network. However, the *opportunistic* transmission mechanism trades throughput and latency for reducing transmission costs.

We present an entirely different concept of utilizing intermediate quality links that are currently ignored by routing protocols. Our approach utilizes long-range bursty links in the network when they temporarily become reliable for transmission.

2.3 Opportunistic Routing

Opportunistic Routing [4] in 802.11 based wireless networks reports a throughput increase of 35% by utilizing long range wireless links. It uses an agreement protocol among the intermediate nodes that receive a batch of packets for prioritizing the intermediate node closest to the destination for forwarding packets. However, it has a relatively high overhead with regard to computational cost, storage, and communication which is not feasible in resource constrained sensor networks. Opportunistic routing operates on a batch of

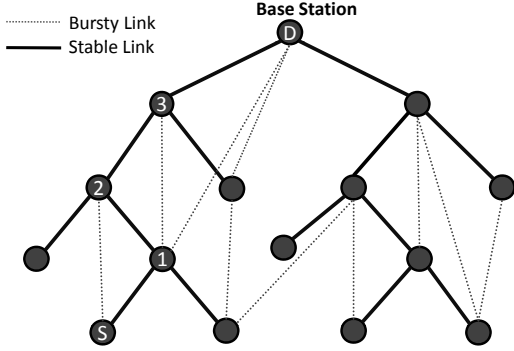


Figure 1. Bursty links provide routing shortcuts that can significantly reduce the hop count and the number of transmissions from source to destination.

packets and strives to reach a delivery threshold of 90% before turning back to traditional routing for delivering the remaining 10% packets. In contrast to opportunistic routing, our primary goal is to reduce the number of transmissions in the network. We apply unicast forwarding and hence the next forwarder of the packet is predetermined. Similarly, our approach rapidly falls back to traditional routing to avoid overshooting links with high loss rates. Our goal is to utilize long-range bursty links to increase routing progress and throughput without introducing significant overhead in terms of computation, storage, and communication.

Overall, our short-term link estimator and its integration with routing protocols is designed according to lessons learned from experimental studies on bursty wireless links, such as [3, 26]. However, our work does not aim at modeling and developing the analytical or experimental understanding of wireless links. Instead, we take a step further and use these experimental models for packet forwarding over bursty links, and hence, enabling better utilization of wireless links.

3 Overview

In this section we present the basic concept of BRE: packet forwarding over bursty links. We also discuss the main goals and challenges of integrating bursty links into the routing process.

3.1 Basic Concept

Typically, routing protocols in WSNs aim to establish a routing tree: Some number of nodes in the network would advertise themselves as base stations, i.e., as tree roots. All other nodes join the tree with ETX as the routing metric. Figure 1 shows an example of such a routing tree rooted at the base station D . A path from source S to the destination D consists of a sub-sequence of immediate parents of each node, for example $S \rightarrow 1 \rightarrow 2 \rightarrow 3 \rightarrow D$. If we consider all links in this path to be 100% reliable, the minimum number of transmissions required by a packet to travel from the source to the destination is four. Now consider a situation in which an intermediate link $S \rightarrow 2$ or $1 \rightarrow D$ has become temporarily reliable. Routing over these links could result in a path sequence $S \rightarrow 2 \rightarrow 3 \rightarrow D$ or $S \rightarrow 1 \rightarrow D$, respectively. Hence, using these links for routing could reduce the total number of transmissions to three in the former and two in the latter case. However, a traditional routing protocol does not

make use of such an opportunity because it uses a long-term link estimate. Hence, this design is intentionally unable to realize short-term changes in the link quality. Similarly, even if these short-term changes are captured, traditional routing schemes adapt slowly to ensure routing stability.

In contrast, our proposed technique takes advantage of the availability of intermediate links. It estimates links on a short-term basis by overhearing packets. In this particular case for example, node-2 overhears the packets addressed to node-1 by source S . After node-2 successfully overheard a certain number of consecutive packets from S , it informs S about the short-term availability of this link. Thereafter, S starts forwarding its packets to node-2 to reduce the number of overall transmissions for a packet to reach its ultimate destination.

The packet overhearing technique employed in STLE benefits from the fact that WSNs typically reveal bursty traffic patterns. Common applications [13, 18, 23, 28, 29] operate as monitoring environment to detect and often track events. Typically, their occurrence results in long bursts of packets. Hence, they represent a major fraction of the overall network traffic although they occur rarely. In such situations, STLE, after overhearing the first few packets over a bursty link, identifies it as short-term available for transmission.

For example, consider a simple sensor-network-based fence monitoring system [29]. During normal conditions, i.e. when there is no intruder breaking into the fence, the network generates very limited or no data traffic. However, as soon as an intruder is detected by the system, large bursts of packets are generated by the distributed event detection algorithm. In such a situation, STLE recognizes bursty links currently reliable for transmission and informs the routing algorithm about the availability of such links.

Our evaluation results in Section 5 show that this technique significantly reduces the number of overall transmissions in the network. Thus, it allows to reduce energy consumption and thereby increase network life time.

3.2 Design Goals and Challenges

Our major design goal is to reduce the number of transmissions in the network and increase routing throughput by utilizing long-range bursty links for packet forwarding. However, we seek to achieve our goal without affecting the reliability and stability of traditional routing. Therefore, our approach of transmitting over links with high loss rates faces four key challenges that influence our design decisions.

First, routing over temporary available links increases the risk of packet loss. Hence, STLE shall accurately predict the periods of reliable transmissions in bursty links. Similarly, it shall provide a backup mechanism when there is no bursty link available for transmission.

Second, only the bursty links which offer good routing progress and do not disrupt the stability of the underlying routing topology shall be identified and reported to routing protocols. Failure to meet this requirement could result in typical routing problems such as loops and network partitioning.

Third, STLE and its associated routing strategy should be lightweight and resource sensitive in terms of computation, storage, and communication. For example, frequently

broadcasting beacons to estimate a link on short-term basis is prohibitive because it would consume significant amounts of energy and bandwidth, the two most critical resources in sensor networks.

Fourth, the short-term link estimation technique should seamlessly integrate with existing routing protocols and link estimators. Its use shall not affect applications and services running on top of routing protocols. However, applications must be amenable to path changes of STLE and not expect a static routing path to the destination.

4 Design

In this section we discuss the design of BRE, which consists of three basic components: (1) STLE, which identifies periods of good transmissions in long-range bursty links by overhearing communication channels, (2) an adaptive routing strategy, which makes use of bursty links identified by STLE for forwarding packets, and (3) the integration of STLE and the adaptive routing strategy with existing routing protocols and link estimators.

4.1 Short Term Link Estimation

The main task of STLE is to identify reliable periods of transmissions in intermediate links that offer better routing progress than long-term stable links. For this purpose, STLE overhears data packets sent by neighboring nodes and records the recent history of success or failure over a link. Based on this recent transmission history, STLE decides whether a link is currently reliable or unreliable for transmission and informs the routing protocol accordingly. Thus, due to its passive design, it estimates the quality of a link without transmitting periodic control packets.

In the following sections, we discuss the algorithm for STLE and the required *thresholds* for the recent transmission history of a link (Section 5.3 presents the experimental evaluation for the settings used in the algorithm).

4.1.1 Algorithm

Before elaborating the algorithmic details of STLE, we define three roles for nodes in the network: a) *sender-node*: the node which is actively sending or forwarding packets b) *parent*: the parent of any sender-node in traditional routing and c) *overhearing-node*: node(s) which can overhear the communication between the sender-node and its parent. A node in the network can assume any or all of these three roles at a time. The STLE algorithm has the following four phases:

Link Reliability:

When overhearing a packet from a sender-node, the overhearing-node infers the success rate - derived from the packet sequence number - of the link with that sender-node. If the loss equals zero, i.e., if the overhearing-node was able to overhear a sufficient number of consecutive packets sent by the sender-node to its parent (see Section 4.1.2), the overhearing-node triggers the next phase of the algorithm. However, if the overhearing-node was unable to overhear a sufficient number of consecutive packets, it drops the oldest packet sequence number for that sender-node from its history and waits for the next packet.

Link Feasibility:

In this phase, the overhearing-node queries the routing protocol for the path-ETX of the packet's destination, i.e., the parent of the sender-node. If the path-ETX of the parent-node is greater than that of the overhearing-node, the overhearing-node declares the bursty link between itself and the sender-node active. Consequently, the active bursty link can offer a better routing progress than the traditional path used by the sender-node. However, if the path-ETX of the parent-node is not known or less than the path-ETX of overhearing-node, the overhearing-node temporarily ignores the sender-node.

Link Announcement:

If the path-ETX of the parent-node is greater than that of the overhearing-node, the overhearing-node informs the sender-node about the active bursty link. It volunteers itself to become the temporary parent of the sender-node as long as this bursty link remains active.

The path-ETX information used by STLE at the overhearing-node can easily be obtained by using the neighborhood information maintained by any traditional routing protocol. We assume that there is a high probability that the original parent of the sender-node is also a neighbor of the overhearing-node. This is because the overhearing-node can listen to the ongoing communication between the sender-node and its parent. Additionally, the link announcement message, sent by the overhearing-node to the sender-node, establishes a simple check to test for link-asymmetry.

Link Unavailability:

At the sender-node, STLE declares a link unavailable for transmission after it fails to receive a number of acknowledgments (see Section 4.1.2) for the data packets sent over the bursty link. In the mean time if there is another bursty link currently active, STLE informs the routing protocol about the availability of this link. However, if no such link is available, STLE backs off and allows the routing protocol to forward packets via the original parent in traditional routing.

4.1.2 Thresholds for STLE

The STLE requires two thresholds for its operation: (1) a threshold to determine after how many successful transmissions, i.e. packets overheard, we define a link temporarily available and (2) after selecting it for routing, a threshold to define how many transmission failures we allow before considering a link temporarily unavailable. Our previous experiments [3] suggest a value of three, i.e. a history of size of three, for the first threshold and one for the second. However, these numbers were derived for a single testbed. As part of our evaluation (see Section 5.3) we repeated these experiments for widespread testbeds such as MoteLab and TWIST to calibrate STLE. Overall, our experimental results suggest the same thresholds as in [3]. Hence, we believe the thresholds are valid in general and not only for a single deployment.

The rationale behind using a history of three successful transmissions is to filter out the bursty links which rarely transmit a batch of packets successfully but would be identified as active for a lower threshold (i.e. $h = 1$ and $h = 2$). It is imperative that our link selection mechanism is not

dominated by such unreliable links. These links would increase the overall transmissions in the network due to frequent packet losses. Similarly, higher values of h (i.e. $h \gg 3$) would result in rarely detecting active bursty links. As a result, our approach would be unable to offer any significant advantage over traditional routing.

The STLE declares a bursty link inactive if only a single data packet is not acknowledged. Experimental results indicate that after a single packet loss over a link of intermediate quality, the probability that the next transmission over an intermediate link will be successful becomes less than 20% [3]. Therefore, we do not risk the retransmission of a packet over the same intermediate link until it is again declared active by STLE. Similarly, the results presented in [26] suggest a back-off time of up to 500 ms before making the next transmission attempt over the same high quality link that has packet losses in bursts. To not incur this latency penalty, our approach instead uses another active bursty link instead of backing off for such a relatively long time.

4.2 An Adaptive Routing Strategy

After discussing the operation of STLE, we now detail a greedy and an adaptive routing strategy based on it. Whenever STLE at a sender-node informs the routing strategy about an active bursty link, the routing strategy makes the overhearing-node its temporary parent and starts forwarding packets to it. However, this information is not propagated by the routing protocol to its descendant nodes, because these short term changes would trigger further parent changes down the tree. Eventually, it might destabilize the routing protocol and result in loops. This is one of the primary reasons why stability prevails over adaptability in today's routing protocols and link estimators. Hence, our routing strategy supplements their design considerations.

The main disadvantage of this approach is that the adaptive routing strategy operates greedily. Although this approach is still effective for enhancing routing progress when compared to traditional routing, it does not promise the use of the optimal path currently available in the network. For example, a node may change its parent based on the recommendation of STLE. However, it is possible that along the traditional path the sender-node remains unaware of the availability of an even better bursty link currently reliable for transmission. Nonetheless, we believe that our approach strikes an efficient trade-off between routing stability and performance adaptability.

The next phase of the adaptive routing strategy deals with falling back to traditional routing, i.e., when STLE declares a bursty link inactive. In this phase, the adaptive routing strategy proceeds as follows:

- It queries STLE for another active bursty link. If such a link is available, the routing strategy starts forwarding packets over it.
- If there is no active bursty link, the adaptive routing strategy will regress to traditional routing until the STLE again finds an active bursty link.

Overall, our approach exploits two types of links that are not utilized by current routing protocols: (1) Short-term stable links (in the order of milliseconds) or bursty links by tak-

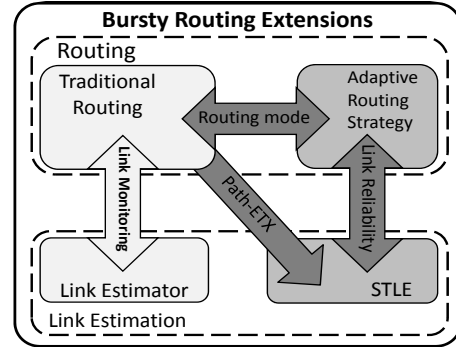


Figure 2. Design of Bursty Routing Extensions

ing into account short-term dynamics which are well beyond the resolution of current link estimators; (2) Relatively long-term stable links (in the order of seconds or even minutes), which offer better routing progress but are not utilized due to slow adaptivity of routing protocols.

4.3 Integration with Routing Protocols

Our goal is to enhance routing performance without affecting the stability and reliability of traditional routing protocols. Therefore, we neither replace the existing link estimators nor alter the stable routing topology maintained by traditional routing protocols. Rather, our approach adds an additional component to the system architecture that assists routing protocols and link estimators in identifying the previously ignored class of bursty links which can enhance routing performance. For this purpose, we define two routing modes, a *bursty mode* and a *traditional mode*. In *bursty mode* packets are forwarded over the active bursty links identified by STLE. Conversely, in *traditional mode* packets are forwarded along the path chosen by the regular routing algorithm. We use the following three interfaces for the integration of STLE and the adaptive routing strategy into traditional routing protocols.

- The first interface is between STLE and the routing protocols. Using this interface, STLE accesses the neighbor table maintained by the routing protocols to enquire the path-ETX of neighboring nodes.
- The second interface is between STLE and the adaptive routing strategy. Using this interface, STLE informs the routing protocol to start or stop forwarding packets over a bursty link when it becomes active or inactive, respectively. STLE examines the availability of a bursty link by monitoring acknowledgements received for each data packet.
- The third interface is between the adaptive routing strategy and routing protocols. This interface is used to switch between different routing modes.

Figure 2 shows the major design components of BRE and their integration with traditional routing protocols. Overall, this design integrates seamlessly into widespread WSN routing protocols such as CTP or BVR.

4.4 Design Challenges

Reliable end-to-end packet transmission and stable network topology are the basic requirements of wireless routing

protocols. One of the major concerns that surfaces with routing over intermediate links is its impact on the routing stability and reliability. Therefore, any approach that attempts to route packets over intermediate quality links needs to alleviate these concerns. Our goal is to benefit from the increased routing progress of specific bursty links. However, we do not want to deteriorate the stability and reliability of wireless routing. In the following sections we address the challenges that stem from packet forwarding over bursty links.

4.4.1 Reliability

To ensure high end-to-end reliability, we have three built-in mechanisms in our approach. First, we eliminate all the bad links that rarely transmit a packet by keeping a recent history of transmission characteristics and waiting for three successful transmissions before declaring a bursty link active. Secondly, our approach employs an aggressive back-off technique to stop transmitting over a bursty link even after a single packet loss (see section 4.2 for details). Both these mechanisms ensure that we do not overshoot a bursty link. Therefore, unlike reactive routing protocols [11, 15, 19] in which route discovery is typically triggered by route break and route timeouts [21], our approach promptly reacts to the changes in link quality.

Finally, as a backup, we use traditional routing and its retransmission mechanisms to deliver the packets that failed over bursty links. The analysis of our experiments in Section 5.4.2 rationalizes that our approach indeed does not affect the reliability of traditional routing.

4.4.2 Stability and Adaptability

Routing stability prevails over performance adaptability in traditional routing protocols [21]. Typically, route evaluation depends on the rate at which beacons are exchanged in traditional routing protocols. However, data is typically exchanged at much higher rates than beacons. Therefore, traditional routing protocols fail to recognize the route quality fluctuations that occur at shorter time scales proportionate to the data exchange rates.

In contrast, always picking the optimal path can itself be detrimental for network performance due to the following reasons; (1) the resulting instability can lead to routing problems such as loops and (2) the overhead associated with active link estimation at shorter time scales is not acceptable for resource constrained sensor networks in terms of energy and bandwidth.

Our adaptive routing strategy finds a suitable trade-off between stability and performance adaptability. Its route evaluation is dependent on the time and the rate at which the data is transmitted and independent of the rate at which beacons are exchanged in the network. Hence, STLE monitors a link by overhearing data packets that, due to the broadcast nature of wireless medium, are received in any case. Thus, STLE ensures link estimation at a high resolution in time with only a small communication overhead. Moreover, local optimizations performed by our adaptive routing strategy in response to the short-term link quality variations are not distributed among other nodes in the network. Therefore, our approach preserves the routing stability by sustaining the routing topology laid down by traditional routing protocols.

As a result, as discussed in the following section, our approach does not introduce routing problems such as loops.

4.4.3 Loops

Loops (or cycles) are a common routing problem in wireless networks which occur due to sudden changes in the routing topology. Loops occur when, due to sudden loss of connectivity to the current parent, a node selects a significantly higher ETX route that also contains a descendant node. A loop is detected when a receiver node finds that its ETX is higher than the ETX of the sender of the packet.

Our adaptive routing approach inherently prevents routing loops. The temporary parent selection mechanism ensures that an overhearing-node is only selected as a new temporary parent if it has a lower path-ETX than the current parent. Additionally, our approach operates locally and does not inform the descendant nodes about the temporary changes made in the parent selection. Hence, our routing strategy, although highly adaptive, does not amplify the looping problem because it preserves the underlying routing topology. Apart from the rare occurrences of loops in traditional routing, we did not observe any additional loops during the evaluation of our approach. Therefore, the loop detection mechanism employed by traditional routing protocols is sufficient, as the integration of BRE does not escalate the occurrence of loops in traditional routing.

5 Evaluation

In this section we evaluate the performance of BRE when compared to CTP. We divide our evaluation in three main sections. After describing the implementation details and experimental setup, we first evaluate the impact of the recent transmission history of a link on the performance of BRE by varying the history-size threshold. This evaluation answers why we use a threshold of 3 in our implementation of STLE. Next we evaluate the impact of BRE on the routing cost and throughput. We give a detailed account of the timing properties of the bursty links used by BRE to enhance routing performance. We conclude our discussion by giving a detailed account of the overhead introduced by STLE.

Our data analysis mainly focuses on routing issues such as transmission costs, delivery reliability, and throughput. Experimental studies, such as [26] and [3], give further insight into the properties of intermediate and bursty links.

5.1 Implementation

We implemented the Bursty Routing Extensions in nesC [9] for TinyOS 2.x. The prototype implementation of BRE consists of STLE and the adaptive routing strategy with CTP, a standard collection protocol for sensor networks shipped with TinyOS. CTP uses the Four Bit Link Estimator (4BLE) as its link estimation component. Although CTP is explicitly designed for relatively low data rates [22], we observed that it is capable of handling high traffic rates as well (i.e. it can deliver a packet every 25 to 30 milliseconds in a multihop network). Moreover, CTP has a very robust retransmission mechanism that ensures high delivery reliability. This property of CTP allows us to thoroughly evaluate the impact of BRE on the reliability of traditional routing protocols. However, STLE and its routing strategy are not




Figure 3. An abstract representation of the MoteLab topology on three different floors. The figure does not show the walls between rooms. The node IDs are only shown for the nodes that were used either as senders or collection roots during our experiments

bound to any specific routing protocol. It can easily be integrated with BVR and other routing strategies that support higher data rates for bandwidth limited systems. Such strategies could, for example, merge multiple data frames into a single link layer packet.

5.2 Experimental Setup

Majority of our experiments were executed on MoteLab, a standard sensor testbed at Harvard University. MoteLab is an indoor deployment of 190 TMoteSky [20] sensor motes on three different floors. However, due to the difficulty of maintaining such a large test-bed, only 142 motes were available to us at maximum. All our experiments had the following common characteristics unless explicitly stated otherwise: (1) Motes transmit at full transmission power i.e. $0dBm$, (2) We use an inter-packet interval of 250 milliseconds (results are presented for different inter-packet intervals as well), (3) We use a neighbor table size of 30 entries for BRE and CTP in all our experiments for consistency. The node density, i.e., the number of neighbors, is much less than 30 in most of the experiments, (4) We use the default $\alpha = 9$ for WMEWMA [31]-based estimation in 4BLE, and 802.15.4 channel 26. The reason behind using $\alpha = 9$ is to ensure a stable routing topology and to avoid loops, (5) Each experimental run lasted for 30 minutes.

To ensure the validity of our MoteLab results, we re-ran our experiments on TWIST, a 100 node TMoteSky-based testbed at TU Berlin. TWIST is a high-density grid-like deployment with an inter-mote spacing of 3 meters. Therefore, to create a reasonably large multihop network, we reduced the transmission power to $-15dBm$ for our experiments on TWIST. The other characteristics are identical to our experiments on MoteLab.

5.3 History Size for STLE

In order to find short-term influences of past packet reception events on future packet reception, we calculate the *conditional* packet reception probability (CRP) for every future packet. The primary reason to calculate the CRP is to

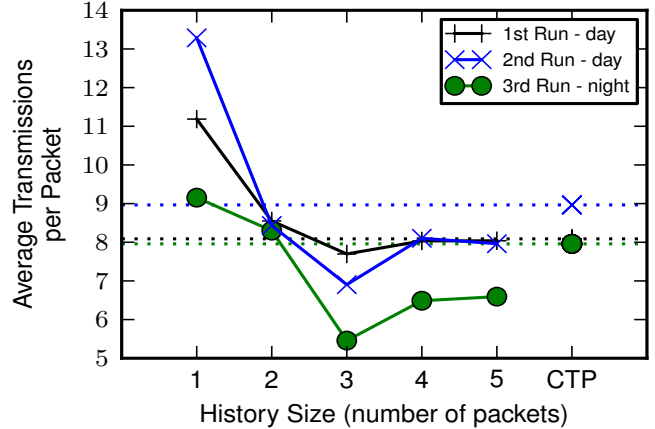


Figure 4. Evaluation of different history size thresholds for STLE on MoteLab. The dotted straight line represent the average of CTP for corresponding experimental runs. $h = 3$ performed the best overall.

calibrate STLE with an optimal *threshold* to determine when an intermediate link should be considered active. The experimental results presented in [3] realize the observation that the CRP increases as the number of preceding successfully received packets over a link increases. Moreover, it concludes that three consecutive successful packet receptions over a link increase the CRP of the next packet to 0.9, and that a history-size threshold of 3 is a sufficient value to detect active bursty links.

In this paper, we grasp these observations to find whether or not this correlation between a history-size threshold of 3 and CRP of 0.9 holds in different testbed environments, such as MoteLab and TWIST. Moreover, we also want to verify the applicability of these results at the routing layer over multiple hops. Figure 4 shows the impact of different history sizes on the performance of BRE for three experimental runs, two during the day and one at night. It can be seen clearly that a threshold of 3 is indeed a sufficient value to predict the future quality of a link as it minimizes the transmission costs when compared to different history sizes as well as a standard routing protocol i.e. CTP. The results in Figure 4 also explain the variations in low-power wireless link qualities over time. Even the delivery cost of a routing protocol like CTP, which restricts communication only to consistently high-quality links, differs by one transmission per packet for back-to-back experimental runs. Nonetheless, the optimality of $h = 3$ holds true for all the three experiments. Therefore, we classify every future packet - the short-term stability of a link - according to the recent *history* of the link, i.e., whether the last h packet receptions were successful or not. In our prototype implementation, we use a threshold of 3, and STLE declares a link as active only after overhearing three consecutive transmissions from a particular sender-node.

To preserve the simplicity of our algorithm for our prototype implementation, we use a static history size. However, we believe that a more perceptive approach could also be useful. For example, one such approach would be to em-

Experimental Class	Intermediate Links %	Forwarders %	Candidates %	Node Density	Potential Neighbors	Candidate Neighbors
Horizontal	33.3	94.8	90.2	15.0	11.2	8.6
Vertical and Diagonal	36.5	93.4	88.4	23.2	14.8	8.5
Nearby	14.2	86.2	79.3	16.3	9.6	4.0

Table 1. MoteLab statistics for experimental parameters defined in Table 2. The statistics for *Intermediate Links*, *Node density*, *Potential Neighbors*, and *Candidate Neighbors* were collected by randomly selecting 10 motes from different locations (i.e. corner, center) in the test-bed. The statistics for *Forwarders* and *Candidates* were collected by running BRE on all the motes (sending a packet every 5 seconds) with a collection root (i.e. mote 183), located at one corner of the network.

ploy a learning phase at the startup for calibrating STLE. Nonetheless, the results in the following section prove the feasibility of this relatively simple approach in principle.

5.4 Performance

In this section we thoroughly evaluate the performance of BRE in terms of transmission cost, throughput, and reliability. Our major performance benchmark is to reduce the number of transmissions in the network by enhancing routing progress. Figure 3 shows a schema of the MoteLab topology² and highlights the motes that were used as senders and receivers in all our experiments. We define four different experimental classes - namely horizontal, vertical, diagonal and nearby - to comprehend different network sizes and topological and physical scenarios (see Table 2). Our mote selection as a source and destination is also based on the these experimental classes.

Before presenting our performance evaluation results, we demonstrate important topology characteristics that describe our analysis and allow for a deep understanding of the results that follow. These parameters are presented in Table 1 and their descriptions are presented in Table 2.

The high percentage of *Forwarders* and *Candidates* in Table 1 shows that a large number of nodes in the network can be utilized in our bursty forwarding approach. Table 1 testifies to the fact that more than 60% (i.e. 11.2 potential neighbors out of 15 neighbors in class *horizontal*) of a node’s immediate neighbors had a better path-ETX than the original parent. Correspondingly, out of these *potential neighbors*, more than 70% (i.e. 8.6 candidate neighbors out of 11.2 potential neighbors in class *horizontal*) could even overhear three consecutive data packets. It means that these neighbors were not selected as a parent only because of a poor long-term quality estimate of their links with the sender, as all of these neighbors did exist in the neighbor table of 30 entries. Algorithms that assess links based on average PRR, like most current approaches, do not use such a link, not even while it is in its good state. The high average of the measured packet loss rate based on broadcast beacons prevents the recognition of good transmission periods in such links.

Another observation is that, with the decrease in the number of intermediate links in the network, the number of *potential neighbors* and *candidate neighbors* also decreases

²This is an abstract representation of the MoteLab topology. A detailed topology and connectivity graphs can be found at <http://www.MoteLab.eecs.harvard.edu>

Name	Description
Horizontal	Source and destination at the opposite ends on the same floor. Only the motes on the same floor were used for this class of experiments (e.g. node-pair 9 → 50).
Diagonal	Source and destination on different floors and on the opposite ends. All the motes in MoteLab were used (e.g. 137 → 50).
Vertical	Source and destination on different floors but on the same end. All the motes in MoteLab were used (e.g. 183 → 50).
Nearby	Source and destination are nearby to each other but surrounded by a high density of nodes. Only 30 to 50 neighboring motes were used (e.g. 153 → 183).
Intermediate Links	The percentage of links in the network with average PRR less than 90%
Forwarders	The percentage of the overhearing-nodes in the network that can overhear a data packet and have a lower path-ETX than the path-ETX of the parent of the sender
Candidates	The percentage of the overhearing-nodes in the network that can overhear three consecutive data packets and have a lower path-ETX than the path-ETX of the parent of the sender
Node Density	Number of neighbors that can overhear a node’s data packet
Potential Neighbors	Number of neighbors that can overhear a node’s data packet and have a lower path-ETX than the path-ETX of its parent.
Candidate Neighbors	Number of neighbors that can overhear three consecutive data packets from a node and have a lower path-ETX than the path-ETX of its parent.

Table 2. Description of experimental classes and parameters presented in Table 1.

(see Table 1 for class *nearby*). Although the *node density* of experimental class *nearby* is higher than class *horizontal*, the class *nearby* has a smaller number of *candidate neighbors*. It means that CTP (based on the link estimates of 4BLE estimator) indeed selected the best neighbor as a parent from the neighbors with high quality links. This information supports the hypothesis in [3, 31] that WMEWMA-based link estimators performs well on good quality links. However, they perform poorly on the links of intermediate quality.

5.4.1 Transmission Cost

We compare the transmission cost of BRE with CTP. Figure 5 shows our results for 16 randomly selected node-pairs as senders and collection roots. To observe the stability of results over time, we repeated our experiments for BRE and CTP three to five times for each of the 16 node-pairs depending upon the difference in number of transmissions. For




Figure 5. Transmission cost reduction and reliability comparison of BRE and CTP. The graph above shows average number of transmissions per packet using BRE and traditional CTP for our experiments on MoteLab. The graph below shows end-to-end packet loss for the same experiments. The bar represents a node pair’s average of five experiments. The error bars represent the highest and the lowest average of the five experiments. The inter-packet interval is 250 ms. For these experiment, the average retransmissions is 8.05% for BRE and 3.5% for CTP. The reduction in the number of transmissions in the case of BRE is mostly due to the reduction in the number of hops.




Figure 6. Average number of transmissions per packet for single experimental runs on TWIST. The error bars in this case represent the standard deviation. The results are similar to the MoteLab experiments.

example, if the average number of transmissions in the case of CTP for a particular node-pair differed by more than one transmission after three experiments, we executed the experiments five times. By reprogramming all the motes involved in an experiment for each experimental run, we enforce CTP to reestablish its routing tree. As a result, we intensively validate our results for a particular node-pair. In most of the cases BRE performs better than CTP, averaging to approximately 19% overall reduction in the transmission costs, i.e., the total number of transmissions from source to destination for single node-pairs.

Although BRE decreases the total number of transmissions in the network by reducing the number of hops, it in-

creases the number of retransmissions when compared to CTP. This is because it risks transmission over links with high loss rates and retransmits all the lost packets via traditional routing. The percentage of retransmissions is 8.05% for BRE and 3.5% for CTP in the experimental results presented in Figure 5.

To see if these results carried over to other networks, we repeated our experiments on TWIST³ using a lower transmission power of $-15dBm$ (see Figure 6). These results are similar and sometimes even better than the results for our experiments on MoteLab. The presented results for an overall of 23 different node-pairs from two different testbeds demonstrate the feasibility of our approach.

There are only a few cases (e.g. node pair 140 → 137) in which CTP is marginally better than BRE. This is due to a simple design tradeoff in our prototype implementation of BRE: For analyzing the precise impact of transmission over intermediate links, currently, we always select an intermediate link without assessing the risk of transmission over such a link. For example, always selecting an intermediate link which has a higher loss rate and only offers a mere 0.1% reduction in transmissions is not always feasible. Frequent failures of transmission over such a link can increase the overall number of transmissions in the network, as depicted in Figure 5 for node-pairs 140 → 137 and 23 → 9.

5.4.2 Reliability

Figure 5 also presents the end-to-end packet loss for our experiments. In most cases, packet loss is negligible. From

³The privacy rules of TWIST did not allow us to show the exact locations and IDs of the node pairs used in our experiments. We used the motes placed on the opposite corners (e.g. south-east and north-west corners) and different floors as senders and collection roots in the grid-like TWIST deployment.




Figure 7. Evaluation results for measured throughput on MoteLab and TWIST. BRE increases routing throughput of traditional routing in most of the cases. The last three bar-pairs show the results for our experiments on TWIST.

these results, it is fair to conclude that BRE does not affect the reliability of the underlying routing protocol and at the same time reduces the number of transmissions in the network. Using the adaptive routing strategy, BRE makes an attempt to forward packets over long-range bursty links. However, when it fails to transmit a packet over a bursty link, it backs off and allows CTP to retransmit the packet over the traditional path.

The only measurable end-to-end packet loss observed in our experiments is for the node-pair $87 \rightarrow 129$ and $87 \rightarrow 67$. We regard these two node-pairs as a sanity check for BRE, as they possess a very lossy path. The average number of hops traversed by each packet for these node pairs is 3 and 4 respectively. However, the average number of transmissions per packet is approximately 8. Therefore, the average link quality is less than 50% in both cases. The upper graph in Figure 5 shows that BRE performs better than CTP even in such lossy scenarios. Similarly, the average end-to-end packet loss for BRE in the case of $87 \rightarrow 129$ is less than in traditional routing. However, as discussed in Section 5.5.2, these two node pairs incur a higher transmission overhead.

Although these node-pairs are surrounded by a large number of motes, as shown in Figure 3 as well as in official MoteLab connectivity maps on the web, we calculated the node density and link qualities for mote 129 to find the exact reasons of this high packet loss. The average PRR for node 129 was less than 40% for all neighbors, and node density was 4 as of February 2009.

5.4.3 Throughput

The two key factors that impact routing throughput in a multihop WSN is the number of retransmissions and the number of hops. Routing throughput can be increased by minimizing the number of retransmissions for a packet to travel from source to destination. Similarly, each hop traversed by a packet also negatively impacts the throughput. The modest computational capability of a sensor node and protocol-specific considerations result in additional delays, such as packet processing requirements and CSMA-backoff waiting time. BRE adapts to both these key factors. Although it slightly increases the number of retransmissions in

the network, the significant reduction in the number of hops contributes to increasing routing throughput (see Figure 7).

CTP is not an ideal candidate for throughput measurements, as it is a reliable routing protocol originally developed for relatively low traffic rates [22]. We still believe that it can provide us with useful hints about the significance of our approach in terms of routing throughput. Furthermore, we use CTP because we wanted to evaluate the maximum throughput without affecting the delivery reliability - a key property of sensor network routing. Our technique for evaluating throughput is to send a packet by calling the *Send* interface of CTP immediately after CTP signals a *SendDone* event for the previous packet. Figure 7 presents our throughput evaluation results⁴ for MoteLab and TWIST. It shows that in most of the cases, due to the reduction in the number of hops, BRE improves the routing throughput - with a maximum improvement of 21%. We expect our approach to be more beneficial if integrated with routing protocols supporting high traffic rates. Moreover, the room for throughput improvement in a bandwidth limited system, like a WSN, is very limited: Langendoen [17] reports a maximum link throughput of 3KB/s for CC2420 without routing in TinyOS. Therefore, in addition to our primary goal of reducing the number of transmission, the throughput increase revealed in Figure 7 is a welcome improvement in a multihop WSN.

Concluding our performance evaluation results, BRE reduces the number of data transmission in the network without affecting delivery reliability. Additionally, by reducing the number of hops for a packet to reach from source to destination, it also enhances routing throughput.

5.4.4 Comparison with a Strawman

In this section we compare BRE with a simple strawman approach - where if a node with lower path-ETX overhears

⁴These results are not comparable for corresponding node pairs in our performance measurement results in Figure 5. The reason is that all our experiments were carried out in a span of 3 months. The MoteLab topology changed significantly during that period. This is also the reason that we had to use different node pairs for throughput evaluations.




Figure 8. Impact of node density on the performance of BRE.




Figure 9. Impact of Table-size on the performance of BRE for Motelab.

a packet, it simply forwards it immediately, without updating any tables. The duplicate packets that arrive along the standard path are later dropped by the overhearing nodes. Comparison with a strawman allows to understand the limits and tradeoffs between the transmission cost and throughput of BRE (see Table 3). We performed our experiments by selecting a single node (node 183) as a root at one corner of Motelab, while other nodes (numbered in Figure 3) sending one at a time a total of 500 packets each. The results clearly show that, while strawman increases the routing throughput by 6%, it worsens the number of transmissions by a factor of 1.8 when compared to BRE.

5.4.5 Node Density and State Maintenance

In this section we analyze how node density and the state of neighboring nodes maintained by BRE impact its performance. Node density positively impacts the performance of BRE as it has more neighboring nodes to choose from. Similarly, higher density increases the probability of finding neighboring nodes with lower path-ETX. This trend is shown in Figure 8. A similar trend can also be seen when comparing different experimental classes presented in Table 1 and the corresponding node-pairs in Figure 5. The node-pairs that belong to high-density experimental class *vertical and diagonal*, such as $137 \rightarrow 37$ and $67 \rightarrow 137$, achieve higher reduction in transmissions.

Finally, we evaluate the impact of table size on the performance of BRE. There are two important algorithmic aspects of BRE that rely on routing table size, (1) the inclusion of neighbors with intermediate link quality in the table and (2) the presence of the original destination of the packet in the routing table of overhearing-node (see section 4.1.1). Figure 9 shows that BRE achieves a very small performance gain for neighbor-table sizes of less than 10 entries

Protocol	Transmissions per packet	Throughput bytes/sec
BRE	4.34	1677
CTP	5.25	1583
Strawman	6.88	1793

Table 3. Summary of the results for BRE and CTP when compared to a strawman. Strawman increases the throughput and the number of transmissions by a factor of 1.06 and 1.8 respectively, when compared to BRE.

on Motelab. The reason is that neighbor-table management policies [7, 31] employed by current link estimators, such as 4BLE, always prefer neighbors with high quality links over the neighbors with intermediate quality links. Therefore, if the number of neighbors with good links is greater than or equal to the table-size, which is the case in Motelab (see Table 1), the chances that BRE will find an entry in the table with lower path-ETX are negligible with existing table-management policies. However, BRE heavily relies on the existence of neighbors with intermediate links in the table. For higher table sizes (i.e. more than 10 entries) these intermediate neighbors automatically get inserted in the table and thus benefit BRE. Similarly, it also increases the probability of finding the original destination of a packet in the table for path-ETX comparison at the overhearing node. Developing table management policies that support BRE by including neighbors with intermediate links into the table is beyond the scope of discussion in this paper, and therefore, we regard it as a future work.

5.5 Intermediate Link Characteristics

After evaluating the performance of BRE, we now analyze the properties of bursty links in more detail. First, we examine the level of correlation between transmission reduction and the number of bursty links used for transmission. Next we present empirical traces from our experiments to investigate the nature and timeliness of intermediate links used for packet forwarding. Finally, we evaluate the impact of different transmission speeds on the performance of BRE.

5.5.1 Transmissions Cost vs. Intermediate Links

In this section, we observe the number of overall packets that were transmitted over one or more bursty links throughout the duration of an experiment. Figure 10 shows that in most of our experiments, more than 50% of the packets were transmitted successfully over one or more bursty links. From these results, it is adequate to conclude that bursty links exist over time in the network and that they are short-term reliable for transmission.

Figure 11 depicts that there is a weak correlation between cost reduction and the number of used bursty links. This fact can also be seen in figure 10 by considering node-pairs $50 \rightarrow 137$ and $23 \rightarrow 9$. The node-pair $50 \rightarrow 137$ achieves 15% reduction in transmission where only 35% packets taking one or more bursty link. Whereas, node-pair $23 \rightarrow 9$ achieves no reduction in transmissions even when 100% of the packets were transmitted over one or more bursty links. The reason for these variations in the results is that, currently, for simplicity and to find the implications of transmission over intermediate links, we do not classify these links and the




Figure 10. Average number of packets transmitted over one or more bursty links vs. reduction in the number of transmissions for the node-pairs as in Figure 5. A large number of packets took one or more bursty links on the path from source to destination in most of the experiments. There is no correlation between the number of packet transmissions over bursty links and the reduction in overall transmissions. For example, in $23 \rightarrow 9$, about 100% of the successful packets took one or more bursty links but did not reduce the number of transmissions in the network. However, in $50 \rightarrow 137$, upto 35% packet transmissions over intermediate links result in 15% reduction in the number of transmissions.




Figure 11. Correlation between the number of packet transmissions over bursty links and the reduction in overall transmissions.

associated overhearing-nodes. As a result, STLE can repeatedly select a bursty link for transmission even if it offers only a slight advantage (e.g 0.1% lower path-ETX) over the link with a traditional parent. However, by limiting the selection of intermediate links only to the ones that offer an adequate cost reduction (such as at least a 10% lower path-ETX), we can improve the correlation between the transmission costs and the number of intermediate links. Similarly, by classifying overhearing-nodes based on their success history, we can eliminate even these rare occurrences of unanticipated results, such as for $23 \rightarrow 9$ and $140 \rightarrow 137$ in Figure 10.

5.5.2 Timeliness

Another property of bursty links that we investigate is timeliness: how often do they occur and for how long are they active. Figure 12 presents empirical traces from our performance evaluation experiments. It shows that bursty links are regularly available over time and are reliable for variable durations. Figure 13 shows the average consecutive packet transmissions over bursty links in each of our experiment. Some of these links are active for only a few milliseconds (e.g $153 \rightarrow 183$), while others for seconds and even minutes

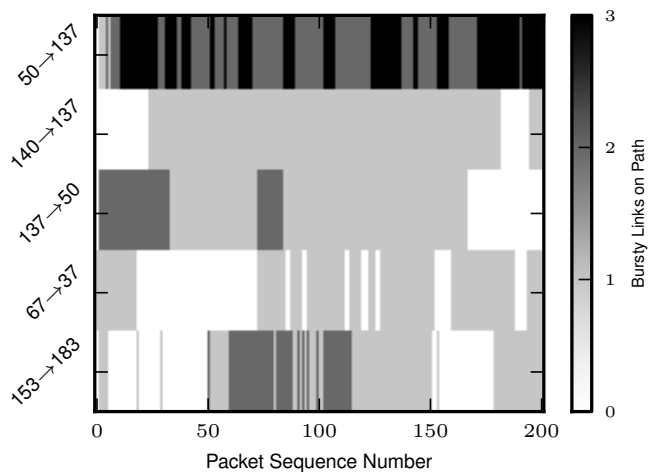


Figure 12. Timeliness of bursty links for 50 second empirical traces for selected node-pairs. The graph shows the variability in the duration for which intermediate links are reliable. Most of the successful packets took one or more bursty links on the path from source to destination. Only the white segments in the graph represent complete packet transmissions on traditional path.

(e.g $140 \rightarrow 37$ in Figure 12). However, due to the slow adaptivity of traditional routing, i.e. CTP, even these relatively long-term reliable links with higher routing progress would not be utilized. Figure 14 shows the cumulative distribution of the burst lengths for all the experimental results presented in section 5.4.1.

5.5.3 Inter-packet Intervals

We investigated the impact of different inter-packet intervals on the performance of BRE. Figure 15 shows that the reduction in the number of transmissions decreases with




Figure 13. Availability of Bursty Links in packet durations. This figure depicts that even relatively long-term (i.e. 750 packet durations) reliable links were also not utilized by CTP. It also shows the limited transmission overhead incurred by BRE.

an increase of the inter-packet interval. This is because sending packets at higher rates over bursty links maintains a strong correlation between their success or failure providence. While by sending packets further apart, the packet loss during a certain measurement period becomes independent [26]. Thus, the correlation between $h = 3$ and $CRP = 0.9$ does not strongly hold for very low transmission rates. It means that at lower transmission speeds, it is less probable that STLE declares a bursty link as active in time. Nonetheless, with inter-packet intervals as high as one second, BRE still offers a 5% improvement of the transmission efficiency compared to traditional routing.

However, as discussed in Section 3.1, we target WSN applications with bursty traffic patterns. Such bursty traffic patterns are typically observed in tracking, monitoring, and surveillance applications. In these applications, the inter-packet interval is expected to be much lower than 1 second during peak traffic times, i.e., the times when the sensor nodes are triggered to track or monitor an activity and report it to the base station.

5.6 Overhead

We divide the overhead introduced by BRE into four different categories namely overhearing, processing, storage, and transmission. The passive overhearing technique that we employ in STLE comes at a cost because a node has to listen to the packets that are not addressed to it. However, due to the broadcast nature of wireless transmission, these packets are always received if the node’s transceiver is in the receive state. State-of-the-art radio chips, such as the CC2420, can be configured to discard all the received packets that are not addressed to a node. Therefore, the overhead associated with overhearing amounts to packet reception and the processing required to deliver a packet from MAC to the link estimator.

The storage overhead of STLE and the adaptive routing strategy is limited to storing the sequence numbers for the

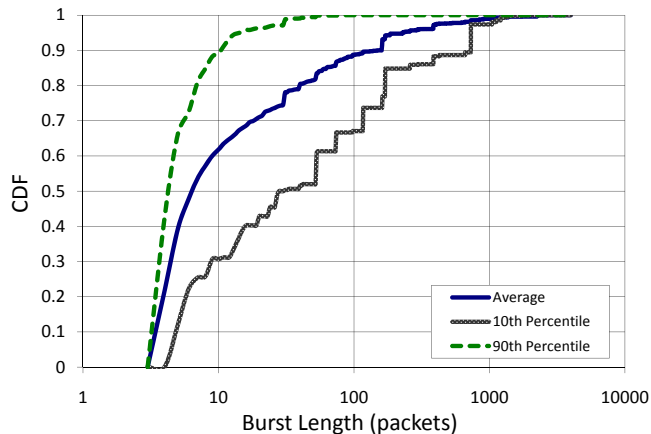


Figure 14. Cumulative distribution of packet bursts for all the experimental results presented in section 5.4.1.

last three packets overheard from the neighboring nodes. Although for consistency in experiments, we use a neighbor table size of 30 entries for BRE and CTP, a smaller table size would also be adequate. Our current implementation of BRE requires 902 bytes of additional code memory and 270 bytes of additional data memory.

Finally, the only transmission overhead introduced by BRE is the announcement message sent by the overhearing-node to the sender-node informing about the temporary availability of a bursty link. There is no retransmission of this message because it also serves the purpose of testing the symmetry of that link. Moreover, as mentioned earlier, even a single successful transmission over a bursty link that reduces one hop would cancel out the overhead introduced by this additional message. However, Figure 14 shows that the burst lengths are much longer for most of our experiments. Considering the fact that BRE can reduce transmission costs by upto 40% and increase routing throughput by upto 20%, we believe that the processing, storage and transmission overhead, presented in this section, is reasonable.

6 Conclusion

In this paper, we presented a simple greedy approach to utilize bursty links of intermediate quality for packet forwarding. Our evaluation results show that, by transmitting over long range intermediate links, the number of transmissions in the network can be reduced.

After evaluating the effectiveness of transmissions over intermediate links, we identify the following aspects as future work: (1) Employing a more perceptive approach for calibrating STLE in different network environments to successfully predict the short-term reliability of a link, (2) Classifying overhearing nodes based on their success history to avoid repeated selection of a node that did not offer significant improvement over the traditional path, (3) Limiting link selection to the ones that offer at least one hop reduction to avoid even the rare occurrence of bad results, (4) Integrating BRE with low-power listening techniques, (5) Extending this work towards 802.11 networks to evaluate whether our approach has a broader relevance in the wireless domain.

The data analysis presented in this paper provides a




Figure 15. Impact of transmission speed on the performance of BRE for node-pair 9 → 50. With the increase in the inter-packet interval, the performance of BRE drops gradually. For the same node pair, the reduction in the number of transmissions drops from 34% at 100 ms to 4.9% at 1s.

greater depth of detail about the extent and applicability of a previously ignored class of links. We believe that the improvement of 19% over traditional routing by transmitting over links with high loss rates is a credible and a realistic result.

7 Acknowledgments

We are very thankful to Geoffrey Werner Challen and Vlado Handziski, who are managing MoteLab and TWIST respectively, for their wonderful contribution to the research community. The amount of data collected for the evaluation presented in this paper would not have been possible without their efforts. Many thanks to Bhaskar Krishnamachari for shepherding this paper and helping us to improve its structure and evaluation quality. This work was supported in part by the UMIC research cluster of RWTH Aachen University and the German Academic Exchange Program (DAAD).

8 References

- [1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. *SIGCOMM Comput. Commun. Rev.*, 34(4):121–132, 2004.
- [2] G. W. Allen, P. Swieskowski, and M. Welsh. Motelab: a wireless sensor network testbed. In *IPSN '05*, 2005.
- [3] A. Becher, O. Landsiedel, G. Kunz, and K. Wehrle. Towards short-term link quality estimation. In *Hot Emnets*, 2008.
- [4] S. Biswas and R. Morris. Exor: opportunistic multi-hop routing for wireless networks. *SIGCOMM Comput. Commun. Rev.*, 35(4):133–144, 2005.
- [5] A. Cerpa, J. L. Wong, L. Kuang, M. Potkonjak, and D. Estrin. Statistical model of lossy links in wireless sensor networks. In *IPSN 05*, 2005.
- [6] 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 *MobiHoc '05*, 2005.
- [7] R. Fonseca, O. Gnawali, K. Jamieson, and P. Levis. Four-bit wireless link estimation. In *HotNets*, 2007.
- [8] R. Fonseca, S. Ratnasamy, J. Zhao, C. T. Ee, D. E. Culler, S. Shenker, and I. Stoica. Beacon vector routing: Scalable point-to-point routing in wireless sensor networks. In *NSDI '05*, 2005.
- [9] D. Gay, P. Levis, R. von Behren, M. Welsh, E. Brewer, and D. Culler. The nesc language: A holistic approach to networked embedded systems. In *ACM SIGPLAN PLDI*, 2003.
- [10] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis. Collection tree protocol. In *SenSys '09*, 2009.
- [11] Govindan and D. Estrin. Directed diffusion: a scalable and robust communication paradigm for sensor networks. In *Proceedings of the sixth annual international conference on Mobile computing and networking*, 2000.
- [12] V. Handziski, A. Köpke, A. Willig, and A. Wolisz. Twist: a scalable and reconfigurable testbed for wireless indoor experiments with sensor networks. In *REALMAN '06*, 2006.
- [13] T. He, S. Krishnamurthy, J. A. Stankovic, T. Abdelzaher, L. Luo, R. Stoleru, T. Yan, L. Gu, J. Hui, and B. Krogh. Energy-efficient surveillance system using wireless sensor networks. In *MobiSys '04*, 2004.
- [14] C. Jiao, L. Schwiebert, and B. Xu. On modeling the packet error statistics in bursty channels. In *IEEE LCN '02*, 2002.
- [15] D. B. Johnson and D. A. Maltz. Dynamic source routing in ad hoc wireless networks. In *Mobile Computing*, pages 153–181. Kluwer Academic Publishers, 1996.
- [16] A. Köpke, A. Willig, and H. Karl. Chaotic maps as parsimonious bit error models of wireless channels. In *INFOCOM*, 2003.
- [17] K. G. Langendoen. Apples, oranges, and testbeds. In *IEEE MASS*, 2006.
- [18] K. Lorincz and M. Welsh. Motetrack: A robust, decentralized approach to rf-based location tracking. In *Springer Personal and Ubiquitous Computing, Special Issue on Location and Context-Awareness*, 2006.
- [19] C. E. Perkins, E. M. Royer, and S. R. Das. Ad hoc on-demand distance vector (aodv) routing. In *2nd IEEE Workshop on Mobile Computing Systems and Applications*, 1999.
- [20] J. Polastre, R. Szewczyk, and D. Culler. Telos: enabling ultra-low power wireless research. In *IPSN '05*, 2005.
- [21] K. Ramachandran, I. Sheriff, E. Belding, and K. Almeroth. *Routing Stability in Static Wireless Mesh Networks*. Springer: Lecture Notes in Computer Science, 2007.
- [22] F. Rodrigo, O. Gnawali, K. Jamieson, S. Kim, P. Levis, and A. Wo. The collection tree protocol. In *TinyOS Enhancement Proposal, TEP 123*, 2006.
- [23] C. Sharp, S. Schaffert, A. Woo, N. Sastry, C. Karlof, S. Sastry, and D. Culler. Design and implementation of a sensor network system for vehicle tracking and autonomous interception. In *EWSN*, 2005.
- [24] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis. Some implications of low power wireless to ip networking. In *HotOS V*, 2006.
- [25] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis. Understanding the causes of packet delivery success and failure in dense wireless sensor networks. In *SenSys '06*, 2006.
- [26] K. Srinivasan, M. A. Kazandjieva, S. Agarwal, and P. Levis. The β -factor: measuring wireless link burstiness. In *SenSys '08*, 2008.
- [27] K. Srinivasan and P. Levis. Rssi is under appreciated. In *Emnets*, 2006.
- [28] M. Welsh, G. W. Allen, K. Lorincz, O. Marcillo, J. Johnson, M. Ruiz, and J. Lees. Sensor networks for high-resolution monitoring of volcanic activity. In *SOSP '05*, 2005.
- [29] G. Wittenburg, K. Terfloth, F. L. Villafuerte, T. Naumowicz, H. Ritter, and J. H. Schiller. Fence monitoring - experimental evaluation of a use case for wireless sensor networks. In *EWSN*, 2007.
- [30] A. Woo and D. Culler. Evaluation of efficient link reliability estimators for low-power wireless networks. Technical Report UCB/CSD-03-1270, EECS Department, University of California, Berkeley, 2003.
- [31] A. Woo, T. Tong, and D. Culler. Taming the underlying challenges of reliable multihop routing in sensor networks. In *SenSys '03*, 2003.