

Comparison of Routing Metrics in 802.11n Wireless Mesh Networks

Ramanujan K Sheshadri, Dimitrios Koutsonikolas
University at Buffalo, SUNY, Buffalo, NY, 14260-2500
Email: {ramanuja, dimitrio}@buffalo.edu

Abstract—We conduct the first experimental study of the performance of link quality-based routing metrics in an 802.11n wireless mesh network (WMN). Link quality-based metrics have been shown to significantly outperform the traditional hopcount metric but they have only been evaluated over legacy 802.11a/b/g radios. The new 802.11n standard introduces a number of enhancements at the MAC and PHY layers (MIMO technology, channel bonding, frame aggregation, short guard interval, and more aggressive modulation and coding schemes) marking the beginning of a new generation of 802.11 radios. Our study in a 21-node indoor 802.11n WMN testbed reveals that the gains of link quality-based metrics over the hopcount metric in legacy 802.11 WMNs do not carry over in 802.11n MIMO WMNs. We analyze the causes of this behavior and make recommendations for the design of new routing metrics in 802.11n WMNs.

I. INTRODUCTION

Wireless mesh networks (WMNs) have experienced a significant development over the past decade, drawing substantial interest in both academia (e.g., [1], [2]) and industry (e.g., [3], [4], [5]). The first use of WMNs [1], [2] was as a low-cost alternative to other last-mile broadband Internet access technologies, such as cable and DSL. However, the advantages of the all-wireless WMN architecture quickly made it an attractive technology for Internet access both in smaller (e.g., all-wireless office [6] or home mesh networking [7]) and in much larger scales (e.g., city-wide WMNs). Regardless of the size, a key design goal in today's WMNs is to provide *high throughput*, comparable to that offered by other broadband technologies, in order to support the ever-growing demands of Internet applications.

A main challenge in WMNs has always been to deal with the poor link quality due to multipath fading [8] and interference [9], both internal (among different hops of the same flow and among different flows) and external (from other 802.11 networks and other devices operating in the same frequency bands). For example, 50% of the operational links in Roofnet had loss rates higher than 30% [8]. To support the high-bandwidth applications, a WMN needs to route traffic over high-throughput paths and avoid highly lossy links. Hence, routing protocol design has always been critical to the performance of WMNs.

In the past few years, researchers have proposed a plethora of link quality-based routing metrics to replace the traditional hopcount (HC) metric in WMNs. These metrics estimate the cost of a path based on the underlying link quality, taking into account factors such as the link loss rates [10], bandwidth [11], or interference [12]. Common to all these techniques is the underlying MAC/PHY layer. These techniques have been designed to work over legacy 802.11a/b/g radios each equipped with a single (omni-directional) antenna,

based on a key observation: the underlying MAC/PHY layers have inherent limitations (limited wireless bandwidth, lossy wireless channels, high MAC/PHY layer overhead) and it is mainly the routing protocol's responsibility to overcome these limitations and provide high-throughput routing paths.

The recent IEEE 802.11n standard [13] marks the beginning of a new generation of 802.11 radios, being the first 802.11 amendment to introduce a PHY layer based on the Multiple-Input Multiple-Output (MIMO) transmission scheme. The MIMO technology, combined with other innovations at the MAC and PHY layers (channel bonding, frame aggregation, short guard interval, and more aggressive modulation and coding schemes) allows 802.11n to provide higher data rates up to 600 Mbps, longer range, and more reliable coverage compared to legacy 802.11a/b/g. Some of the major players in the WMN industry [5], [4], [3] have already introduced 802.11n WMN routers in the market. In a recent study [14] involving 110 commercially deployed WMNs, 33 of them used 802.11n routers.

The MAC/PHY layer enhancements introduced by the new 802.11n standard drastically change the assumptions about the underlying MAC/PHY layers. The use of multiple antennas makes transmissions more robust to channel errors through the spatial diversity MIMO mode or increases the transmission rates through the MIMO spatial multiplexing mode. The latter, combined with aggressive modulation and coding schemes and channel bonding, implies that link bandwidth is no longer a bottleneck (indeed, the highest 802.11n bitrates already offer speeds much faster than typical cable/DSL). Furthermore, frame aggregation, block acknowledgments, and short guard interval significantly reduce the MAC/PHY overhead. Given all these high-throughput enhancements at the underlying MAC/PHY layers, one could argue that intelligent routing is no longer needed for the next generation of WMNs.

Nevertheless, it is not clear whether the large per-link throughput gains of 802.11n carry over in multihop topologies. The main challenge of dealing with the high loss rates prevalent in the multihop WMN environment still remains [14], [15]. The spatial diversity MIMO mode improves the robustness of weak links at the cost of significantly reduced throughput. On the other hand, the spatial multiplexing mode, combined with dense modulation and coding schemes and/or channel bonding, can result in significantly higher loss rates compared to 802.11a/b/g. A recent experimental study [15] showed that the loss rates at the highest spatial multiplexing bitrates can be very high even in the absence of interference. The situation is expected to be worse in WMNs due to multihop communication.

Furthermore, it is not clear whether routing metrics which were proved successful over legacy 802.11 radios will still work well over the new generation of MIMO radios. Link quality-based routing metrics have traditionally relied on probing to measure channel quality. With a large increase in the number of available bitrates along with the newly added features at the MAC/PHY layers, it is not clear whether probing remains an accurate and, at the same time, low-overhead technique for assessing link quality in 802.11n WMNs.

In this paper, we conduct what we believe to be the first experimental study of the two most popular routing metrics, ETX [10] and ETT [11], [1], along with the traditional HC metric, in a 21-node indoor 802.11n WMN testbed. Our study reveals that the throughput gains of link quality-based routing metrics over the HC metric in 802.11a/b/g WMNs do not carry over in 802.11n WMNs. Even worse, the large per-link throughput gains of 802.11n over 802.11a/b/g do not translate into multihop throughput gains. As a striking example, in an 802.11n WMN with frame aggregation enabled and supported bitrates up to 130 Mbps, the state-of-the-art ETT metric (taking into account both link loss rates and link bandwidths) fails to find better paths than ETX (taking into account only loss rates) and its median throughput is 31% lower than in an 802.11g WMN with supported bitrates only up to 54 Mbps.

Furthermore, we investigate the causes for the poor performance of link quality-based routing metrics in 802.11n WMNs. We find that popular probing techniques for link loss rate and link bandwidth estimation yield poor accuracy in 802.11n WMNs due to the new features introduced at the underlying MAC/PHY layers. We also observe that these metrics show a tendency towards short paths, consisting of long links, in 802.11n WMNs, similar to the behavior of the HC metric in legacy 802.11a/b/g WMNs. We end with a few recommendations for the design of new routing metrics over 802.11n WMNs.

The rest of the paper is organized as follows. Section II provides a brief description of the routing metrics we consider in this study, and a brief overview of the main features of 802.11n. Section III describes our testbed. In Section IV, we present our experimental study. In Section V, we discuss challenges in the design of link quality-based metrics in 802.11n WMNs and make a few recommendations based on our findings in Section IV. Section VI discusses the related work. Finally, Section VII concludes the paper.

II. BACKGROUND

In this section, we briefly review the three routing metrics we consider in our study (II-A) and the new features introduced by 802.11n (II-B).

A. Link quality-based routing metrics in WMNs

HC This metric provides minimum hopcount routing. Over the past few years, several studies have shown the disadvantages of this metric in WMNs compared to link quality-based metrics [10], [16], [11].

ETX This metric estimates the expected number of transmissions required to send a unicast packet over a link, including

MAC layer retransmissions. To calculate the ETX of a link, each node measures the probability that a packet successfully reaches the receiver, denoted as d_f , and the probability that an ACK is successfully received by the sender, denoted as d_r , using broadcast probe packets (typically one every 1 second). The ETX value of the link is given by $ETX = \frac{1}{d_f \cdot d_r}$.

ETT This metric estimates the expected time required to send a unicast packet over a link, including MAC layer retransmissions. It is an enhancement to ETX, taking into account both the loss rate and the bandwidth of the link. There are two ways to calculate the ETT of a link [11], [1]. In [11], ETT is calculated as $ETT = ETX \times \frac{S}{B}$, where S is the packet size and B is the link bandwidth, which is estimated using a technique similar to Packet Pair [17]. In [1], each node periodically measures the link delivery probabilities at each available bitrate and predicts that a link's highest-throughput bitrate is the one with the maximum product of delivery probability and bitrate. The ETT for a link is then calculated as the expected transmission time of a 1500-byte packet at that link's highest-throughput bitrate.

B. 802.11n enhancements

Multiple-input Multiple-output (MIMO) The 802.11n standard employs multiple antennas to take advantage of multipath propagation in the RF environment. MIMO links can operate in two different modes: Spatial Division Multiplexing (SDM) and Spatial Diversity or Space Time Block Coding (STBC). The SDM mode utilizes multiple independent data streams transmitted through different antennas over independent spatial paths in order to increase the PHY data rate. Multiplexing two spatial streams onto a single channel effectively doubles the PHY transmission rate. In contrast to SDM, the STBC mode sends a single stream redundantly over multiple antennas. By combining different copies of the signal, arriving over different paths with independent fading characteristics, the receiver has a better chance of accurately decoding the original signal. Hence, STBC provides higher reliability and extended transmission range compared to conventional SISO links.

Short Guard Interval (GI) The GI is the time between transmitted OFDM symbols and is necessary to offset the effects of multipath which would otherwise cause Inter-Symbol Interference. Legacy 802.11a/b/g use a GI of 800 ns, but the 802.11n standard allows for a short GI of 400 ns, increasing the PHY transmission rate by up to 11% while maintaining sufficient symbol separation for dealing with multipath.

Channel bonding Legacy 802.11a/b/g devices operate on 20MHz channels. In contrast, 802.11n devices can use a channel width of either 20MHz or 40MHz. The latter case is known as channel bonding and theoretically doubles the supported data rate.

Supported bitrates 802.11n devices can use a variety of modulation schemes and coding rates. They can also vary the number of spatial streams, the GI, and the channel width. The combination of all these factors results in a large set of supported transmission bitrates, ranging from a minimum of 6.5Mbps up to a (theoretical) maximum of 600 Mbps. The

standard defines *Modulation and Coding Schemes (MCS)* – an integer number assigned to each permutation of modulation, coding rate, GI, channel width, and number of spatial streams, with 8 different bitrates available for each combination of GI, channel width, and number of spatial streams.

Frame Aggregation/Block Acknowledgment (FA/BA)

802.11n allows transmitters to send multiple back-to-back frames at each MAC access opportunity, thus significantly reducing the MAC/PHY overhead compared to 802.11a/b/g. 802.11n supports two types of FA: A-MSDU and A-MPDU. With A-MSDU, multiple higher layer packets (MSDUs) are combined to form an aggregate MAC frame, which contains one MAC header and one checksum. With A-MPDU, multiple MAC frames (MPDUs), are combined to form an aggregate MAC frame with each MPDU carrying its own MAC header and checksum. The use of individual checksums allows the receiver to decode each frame individually and ask for re-transmissions of only those frames received in error, through a BA mechanism which allows a receiver to confirm reception of multiple data frames with a single BA frame.

III. TESTBED

Our study is conducted on UBMesh [18] (Figure 1), an experimental 802.11a/b/g/n 21-node WMN testbed deployed on the 3rd floor of an academic building at the University at Buffalo. Each node is a desktop PC with a Pentium Dual Core E5800 3.2GHz processor and 2GB RAM, running Ubuntu Linux 10.10. Each node has a Ralink RT2860 802.11a/b/g/n mini PCI card operating in ad hoc mode and attached to three 3dBi rubber duck omnidirectional antennas, through low loss pigtailed. The open-source Ralink RT2860 (v2.4.0.0) driver is used to enable the wireless cards. The Ralink RT2860 chipset implements all the available 802.11n features, supports 2x3 MIMO operation, and provides bitrates up to 300 Mbps. Our cards operate in the 2.4 GHz band with a 20 MHz channel width, as the ad hoc mode implementation in the RT2860 driver does not work properly in 5 GHz or with channel bonding. As a result, the supported bitrates are limited between 6.5 Mbps (MCS 0) and 130 Mbps (MCS 7), roughly 2.5 times higher than the maximum 802.11a/g bitrate (54 Mbps). All the experiments were conducted at night to minimize interference from other wireless devices.

We used the OLSR protocol [19] and its publicly available implementations of HC, ETX, and ETT [20], [21]. The ETT implementation is similar to that in [11], i.e., measures link bandwidth using a technique based on Packet Pair and uses the same estimate of loss rate for all available bitrates. The Ralink rate adaptation (RA) algorithm was used in all the experiments.

IV. RESULTS

In this section we discuss the results of our experiments.

A. Overall comparison

Methodology We conducted our experiments for the results in IV-A, IV-B in two phases for 802.11n: with FA/BA turned off and on, respectively. We also repeated the experiments for 802.11g. We followed the methodology of [10]. We consider

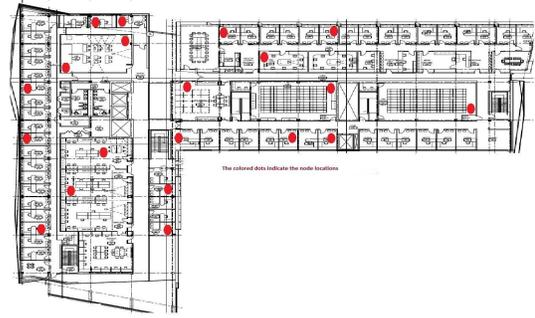


Fig. 1. Floorplan of the UBMesh testbed deployment.

200 node pairs randomly selected from the $21 \times 20 = 420$ ordered node pairs in the testbed. In each phase, we considered each of the 200 node pairs in turn. For each pair, one experiment was run for each metric. At the start of each experiment, we let OLSR run for 120 seconds to stabilize. Then the sending node sent 1500-byte UDP packets at the maximum possible speed to the destination for 10 seconds using Iperf [22]. Differently from in [10] but similarly to in [11], [16], we did not stop OLSR after the 120-second warm-up period. Hence, it is possible that routes change during some of the UDP transfers.

Overall Result Figures 2(a), 2(b), 2(c) plot the Cumulative Distribution Functions (CDFs) of the UDP throughputs of 200 randomly selected node pairs with each metric over 802.11g, 802.11n without FA/BA (802.11n_noFABA), and 802.11n with FA/BA (802.11n_FABA), respectively. As expected, in an 802.11g network (Figure 2(a)), ETX improves the median throughput over HC by 387% (2.53 Mbps vs. 0.52 Mbps), and ETT improves the median throughput over ETX by 33% (3.37 Mbps vs 2.53 Mbps). These results are similar to those reported in [16], [11].

On the other hand, the results are very different in the case of 802.11n (Figures 2(b), 2(c)). Without FA/BA, the median gain of ETX over HC reduces to 27% (3.98 Mbps vs. 3.14 Mbps) and the gain of ETT over ETX reduces to only 8% (4.31 Mbps vs. 3.98 Mbps). With FA/BA enabled, the median gain of ETX over HC further decreases to 12% (2.58 Mbps vs 2.31 Mbps), and the median ETT throughput is slightly lower (2.57 Mbps) than the median ETX throughput. In other words, in a network with bitrates varying from 6.5 Mbps up to 130 Mbps, a metric that takes link bandwidth into account fails to find better paths than a metric that ignores link bandwidth and only uses the loss rate at the lowest bitrate as an estimate of link quality. In fact, the situation becomes worse if we exclude 1-hop paths: ETT performs similar to ETX over 802.11n_noFABA and 11% lower than ETX over 802.11n_FABA.

Note that the 95-th percentile with all three metrics is much larger over 802.11n_FABA (55 Mbps over 802.11n_FABA, 19 Mbps over 802.11n_noFABA, and 30 Mbps over 802.11g). These throughput values correspond to single-hop paths over strong links. For such high quality links, the high SDM bitrates of 802.11n combined with FA/BA can indeed offer large throughput gains over legacy 802.11g. In fact, the maximum

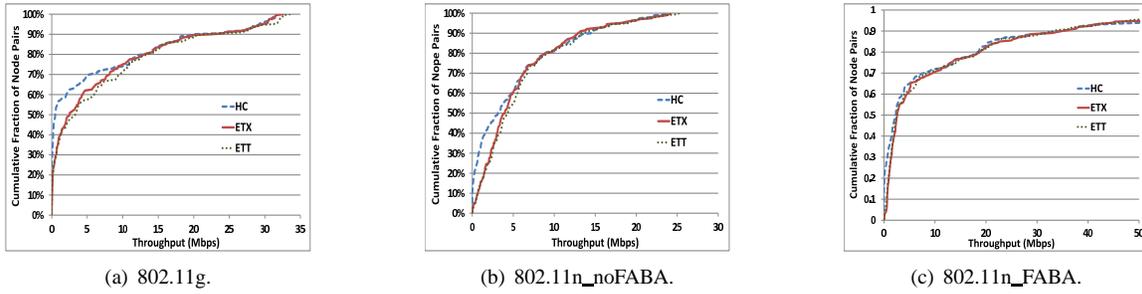


Fig. 2. CDFs of UDP throughputs of 200 randomly selected node pairs with HC, ETX, ETT over 802.11g, 802.11n_noFABA, and 802.11n_FABA.

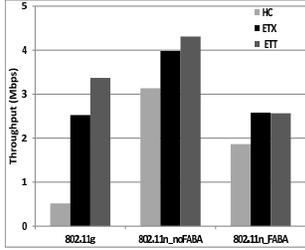


Fig. 3. Median throughput over 802.11g, 802.11n_noFABA, and 802.11n_FABA.

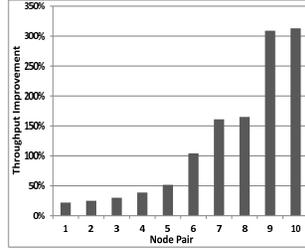


Fig. 4. Throughput improvement of “best” static path over ETT for 10 node pairs over 802.11n_FABA.

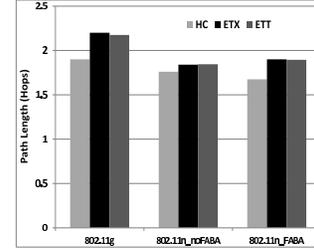


Fig. 5. Average path length over 802.11g, 802.11n_noFABA, and 802.11n_FABA.

throughput we measured over 802.11n_FABA was 90.5 Mbps. **Summary** Figure 3 summarizes the overall result, plotting the median throughput among the 200 node pairs with each of the three metrics when the underlying MAC protocol is 802.11g and 802.11n. We make two key observations: (i) The large throughput gains of link quality-based routing metrics in legacy 802.11g WMNs do not carry over in 802.11n WMNs. In particular, the state-of-the-art ETT metric offers minimal or no improvements over ETX. (ii) Even worse, the large *per-link* throughput gains of 802.11n over 802.11g do not carry over in *multihop* topologies. The median throughput over 802.11n is less than 5Mbps without FA/BA, and even more surprisingly, less than 3 Mbps with FA/BA, with all 3 metrics. Even though the 802.11n enhancements still offer a significant throughput improvement (6x/3.6x with FAB A off/on) with HC, the gain with ETX is limited to only 57% without FA/BA and only 2% with FA/BA. Moreover, the state-of-the-art ETT performs 31% better over 802.11g than over 802.11n_FABA.

The low multihop throughputs over 802.11n, regardless of the routing metric used, could be, in part, a result of poor interactions between the routing protocol and the underlying RA algorithm. RA in 802.11n networks is an active research topic [24], [25] and it is known that several legacy RA algorithms perform poorly over 802.11n links [24]. We plan to investigate this interaction as part of our future work. In this work, we focus on the performance of link-quality-based routing metrics in an 802.11n WMN and their limited gain over HC.

Can we expect higher gains with ETX/ETT in an 802.11n WMN? In Figure 4, we compare the ETT throughput against the throughput of the best statically configured path for 10 randomly selected node pairs over 802.11n_FABA. The “best” path was found by sending data along several potential best paths, one at a time, and selecting the path with the highest

throughput. We observe that the best path offers a 22-313% throughput improvement over ETT. This result suggests that there is a large margin for improvement, however, ETT fails to discover the best paths. In the remaining of the paper, we investigate the reasons for the limited gain of link quality-based routing metrics in 802.11n WMNs, especially when FA/BA is enabled.

B. Impact of path length

Figure 5 plots the average path length with each metric. As expected, HC chooses shorter paths than the other two metrics over both 802.11g and 802.11n. However, the differences in path length are not as pronounced as in previous studies (e.g., [16]).

In Figures 6(a)-6(i), we examine the relation between the path length and the UDP throughput in more detail. These figures plot the throughput of each source-destination pair against its average path length (calculated using 10 ping packets and rounded up to the nearest integer) with each of the three metrics over 802.11g, 802.11n_noFABA, and 802.11n_FABA. In Figures 6(a)-6(c), we observe that ETX and ETT do a good job in avoiding poor 1-hop paths over 802.11g, as expected. The number of 1-hop paths with throughput lower than 5 Mbps is 28 with HC, 8 with ETX, and 7 with ETT. The situation changes over 802.11n (Figures 6(d)-6(i)). The number of 1-hop paths with throughput lower than 5 Mbps with HC, ETX, and ETT is 24, 18, and 19, respectively, over 802.11n_noFABA, and 29, 14, and 15, respectively, over 802.11n_FABA; ETX and ETT often choose poor 1-hop paths similar to HC. We also observe that throughput drops rapidly as the path length increases, with all 3 metrics, especially over 802.11n_FABA.

Another observation from these figures is that the number of 1-hop paths increases for ETX and ETT over 802.11n_noFABA and for all three metrics over 802.11n_FABA compared to 802.11g. The number of 1-hop

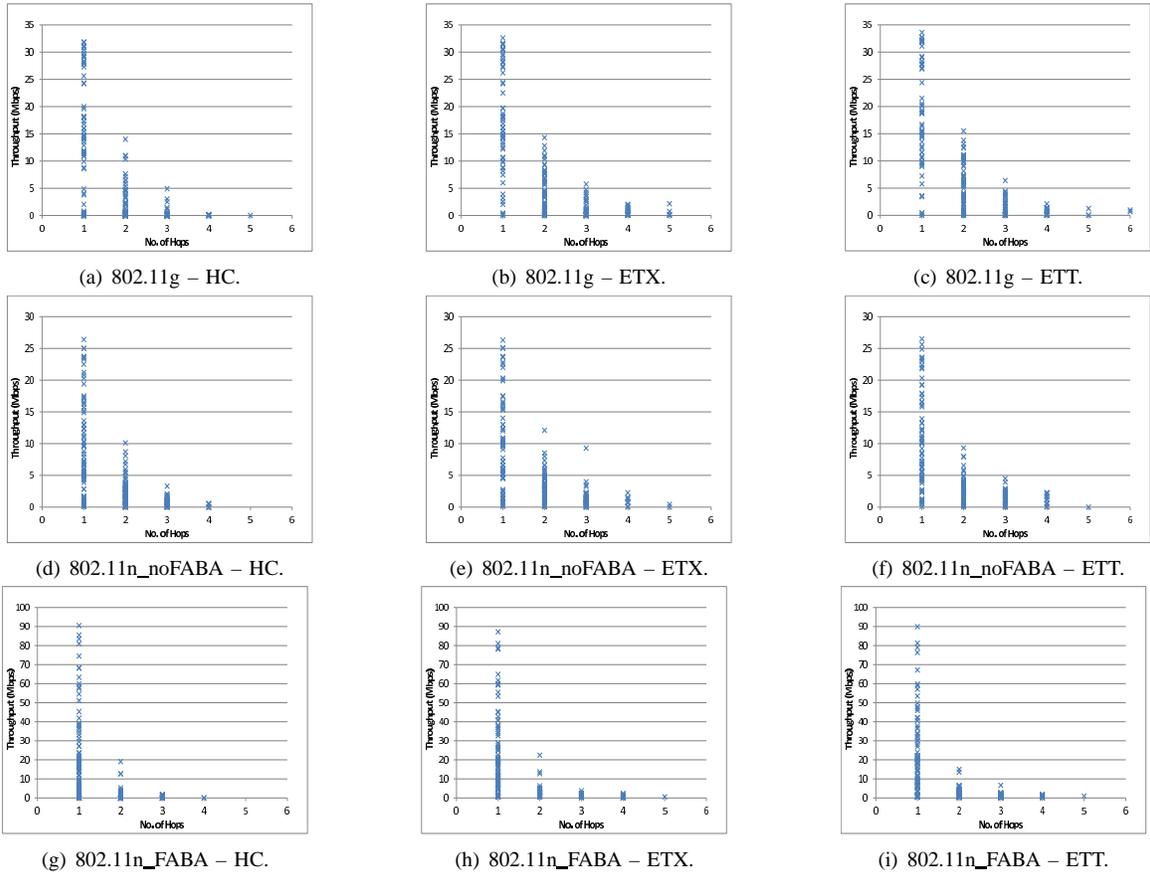


Fig. 6. Throughput as a function of path length.

paths with HC, ETX, and ETT, respectively, is 79, 59, 59 over 802.11g, 77, 71, 70 over 802.11n_noFABA, and 95, 83, 84 over 802.11n_FABA. These results suggest that ETX and ETT prefer paths consisting of long STBC links (DIV-RANGE links in [23]) instead of high-throughput SDM links (MUX links in [23]). One reason for this choice could be that both metrics send probes for loss rate estimation only at the lowest STBC bitrate, resulting in many neighbors per node. A metric taking only loss rate into account or failing to accurately estimate link bandwidth will randomly select among many DIV-RANGE links, resulting in low throughput paths.

In Figures 7(a)-7(i), we examine the impact of path length difference on the throughput improvement between each pair of metrics (ETX vs. HC, ETT vs. HC, and ETT vs. ETX). These figures plot the path length difference between two metrics against the throughput difference for each node pair. We only consider those node pairs for which the path length was longer than 1 hop with at least one metric.

In Figures 7(a), 7(b), we observe that, as expected, ETX/ETT select paths of length larger (typically by 1 hop) than or equal to the length of paths selected by HC over 802.11g. For most node pairs, this selection results in higher throughput compared to HC. In contrast, in Figure 7(c), we observe that most of the paths chosen by ETT and ETX are of the same length. We also observe that ETT improves the throughput over ETX for only 47% of the node pairs. However,

the throughput gains of ETT over ETX are much higher (up to 12 Mbps) compared to the gains of ETX over ETT (at most 3.5 Mbps), resulting in the higher median throughput for ETT (Figure 3).

In Figures 7(d), 7(e), and 7(g), 7(h), we observe that ETX/ETT also select paths of length larger than or equal to the length of the paths selected by HC over 802.11n. However, the throughput improvements with ETX/ETT over HC are lower over 802.11n than over 802.11g. Another difference is that the largest positive throughput differences correspond to path length difference of 0 hops over 802.11n but 1 hop over 802.11g. When we compare ETT vs. ETX in Figures 7(f), 7(i), we observe that again ETX gives higher throughput than ETT for a large fraction of node pairs, similar to over 802.11g. However, the throughput gains of ETT over ETX over 802.11n are similar to the gains of ETX over ETT (in both cases lower than 5 Mbps). Overall, these figures explain the result in Section IV-A. In an 802.11n WMN, ETT and ETX choose mostly paths of equal length and of similar quality for the same node pair. ETT fails to discover better paths than ETX.

C. Throughput variability

In this section, we evaluate the impact of routing metrics on throughput variability. We carried out the following experiment, following the methodology in [16]. We selected 4 nodes on the periphery of the testbed. Each of these nodes carried out a 10 second UDP transfer to the remaining 3 nodes. Only

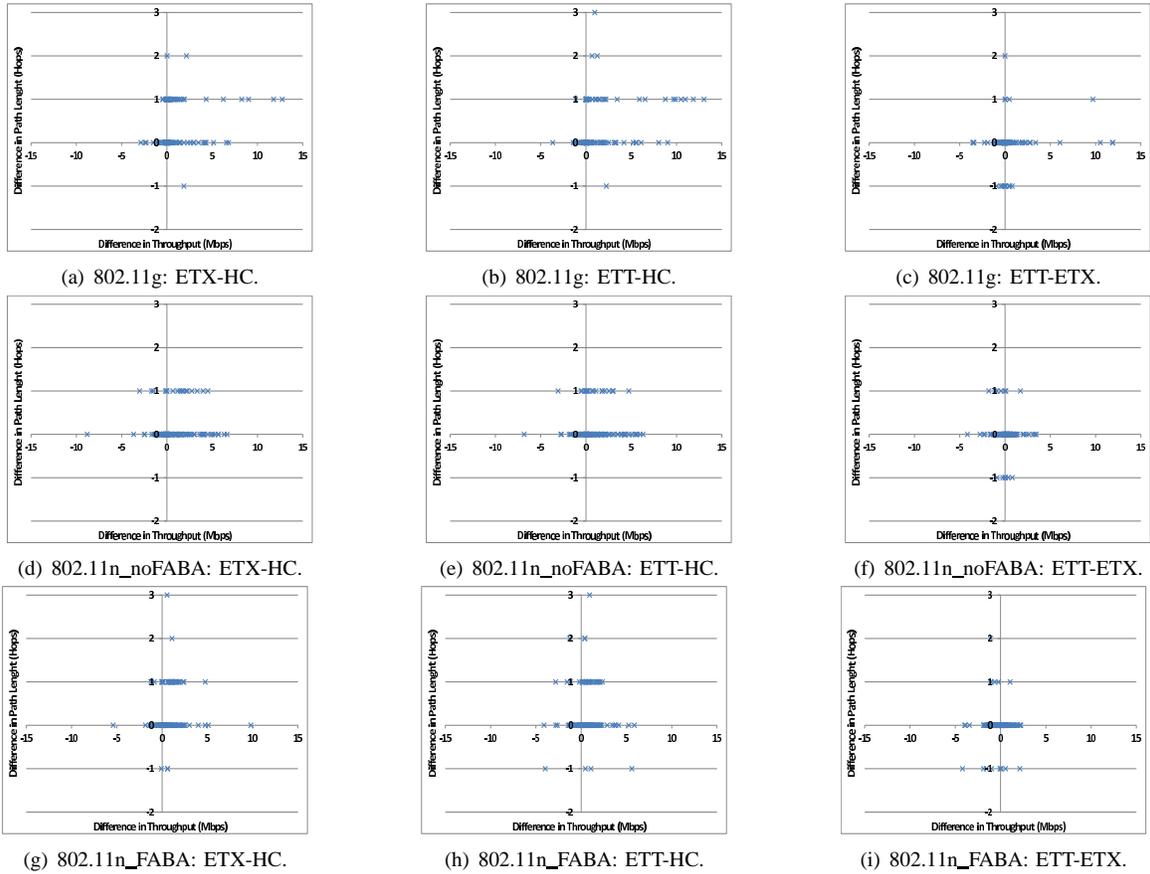


Fig. 7. Scatterplots of path length difference against throughput difference for all node pairs with path lengths longer than 1 hop with at least one metric.

one node was active at any time. We repeated the experiment 10 times for a total of $4 \times 3 \times 10 = 120$ UDP transfers.

Figures 8(a), 8(e) plot the median throughput with each of the three metrics when the underlying MAC protocol is 802.11g and 802.11n_FA. In Figure 8(a), we observe again that ETT offers a significant throughput improvement of 75% over ETX in an 802.11g WMN; this gain is actually higher than the 33% gain observed in the 200-pair experiment (Figure 3). On the other hand, the gain of ETX over HC is only 6%. In Figure 8(e), we observe that the gain of ETX over HC in an 802.11n WMN increases compared to the value in Figure 3 (65% vs 38%). However, the gain of ETT over ETX remains still very low (only 6%), although non-negligible. Moreover, the values with all 3 metrics remain lower than 2 Mbps, similar to those in 802.11g. Again, 802.11n with FA/BA does not improve the multihop throughput compared to 802.11g.

In Figures 8(b)-8(d), 8(f)-8(h), we plot the coefficient of variation (CoV) of the throughputs obtained with each pair of metrics (ETX-HC, ETT-HC, ETT-ETX) over 802.11g and 802.11n, respectively. CoV is defined as the standard deviation divided by the mean and is used as a measure of variability. Each point on these graphs represents the CoV of the throughput of the 10 UDP transfers between the same node pair.

The CoV values of HC are in general much higher than the values of ETX/ETT over 802.11g, (Figures 8(b), 8(c)), which agrees with [16]; link quality-based routing metrics find more

stable paths than HC in an 802.11g WMN. Between the two link quality-based metrics, ETT finds more stable paths than ETX (Figure 8(d)).

In Figures 8(f)-8(h), we observe that the CoV values are much lower over 802.11n than over 802.11g for all three metrics; 802.11n offers higher path stability compared to 802.11g. However, the higher stability in an 802.11n WMN does not translate into throughput gains over an 802.11g WMN, as we saw in Figures 8(a), 8(e). This result again implies that the paths chosen with all three metrics mainly consist of long-range STBC rather than high-throughput SDM links. We also observe that ETX and ETT still find more stable paths than HC, as in an 802.11g WMN. However, unlike in an 802.11g WMN, ETX and ETT offer similar path stability in an 802.11n WMN.

In IV-D, IV-E, IV-F, we try to understand the reasons for the poor performance of ETX and, in particular, ETT in an 802.11n WMN. We discuss the impact of three factors: routing protocol overhead, loss rate estimation, and link bandwidth estimation.

D. OLSR overhead

To evaluate the impact of OLSR overhead, we followed the methodology in [16]. We placed five testbed nodes in close proximity of each other and fixed the bitrate at MCS 15 (SDM, 130 Mbps). We set up static routes between these nodes to

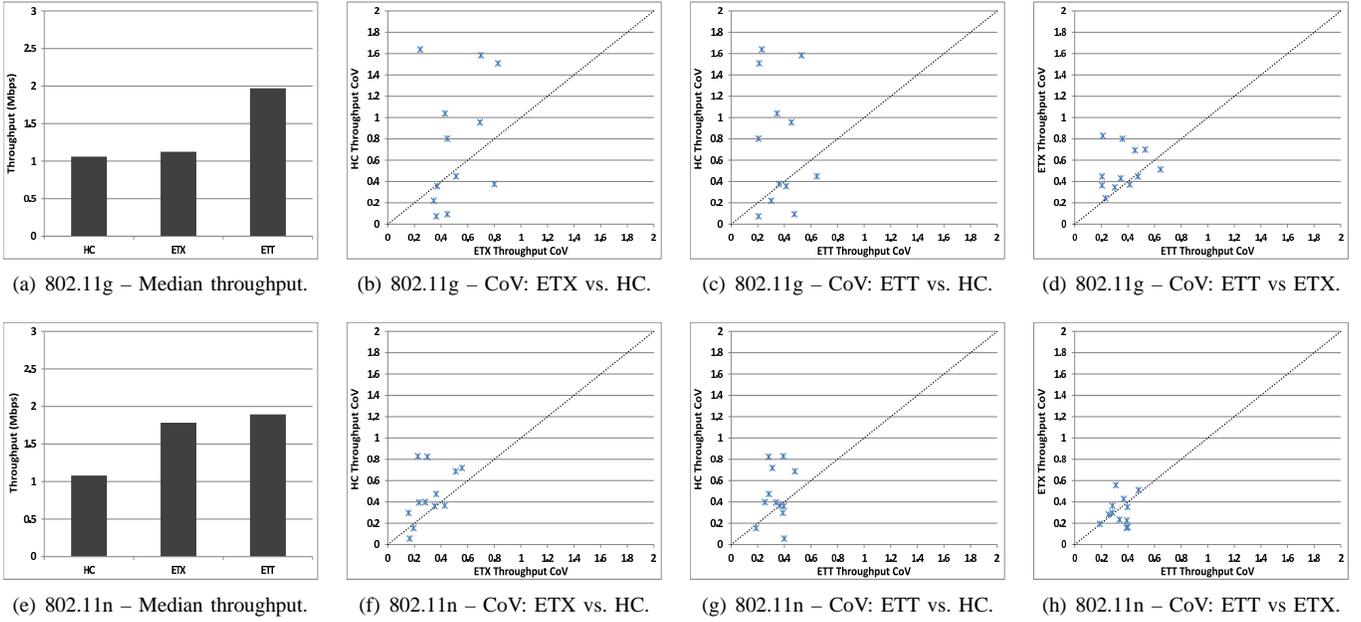


Fig. 8. Throughput variability in an 802.11g and 802.11n WMN with FA/BA.

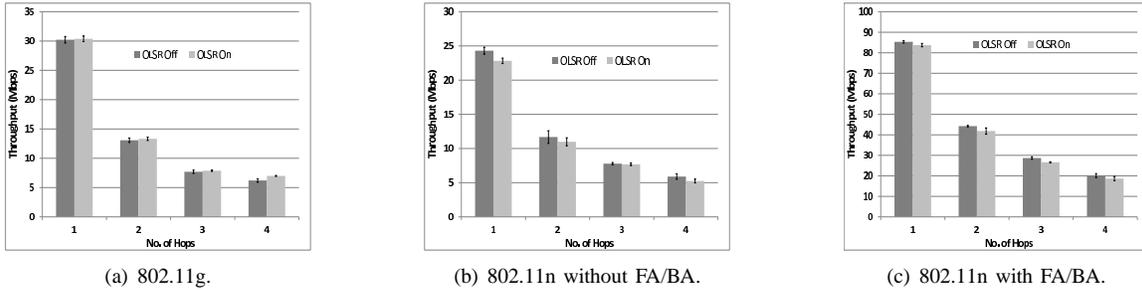


Fig. 9. OLSR overhead.

form a chain topology and measured the UDP throughput of a 1-hop, 2-hop, 3-hop, and 4-hop flow. Then, we started OLSR with ETT on all 21 nodes of our testbed, but we used the same static routes on these 5 nodes overriding the OLSR routes. OLSR continued to send unicast/broadcast probing and routing control packets, but routes discovered through this process were ignored. We measured again the UDP throughput on the same 1-, 2-, 3-, and 4-hop paths. The result of this experiment is shown in Figure 9. Each bar represents the average of 5 measurements.

In Figure 9, we observe that the OLSR overhead is negligible over 802.11g and relatively small over 802.11n. The largest throughput reductions are observed for a 4-hop flow over 802.11n_noFABA and 802.11n_FABA (12% and 8%, respectively). We conclude that the OLSR overhead is not responsible for the low performance of link quality-based routing metrics in an 802.11n WMN.

E. Loss rates

Figure 10 plots the loss rates of the 165 existing links of UBMesh (i.e., links with non-zero loss rate for at least one bitrate) for 3 STBC bitrates (MCS 0, 3, and 7). The links are sorted in increasing order with respect to the loss rates at MCS 0. The loss rates were measured by having each node

in turn broadcast 1500-byte probe packets at each bitrate at a rate of 100 packets per second for 12 seconds. We observe that several links of very good or excellent quality at MCS 0 exhibit extremely high loss rates (often close to 100%) at MCS 3 and MCS 7. In particular, the loss rate varies by more than 21 percentage points for 38 links and by more than 70 percentage points for 23 links when the bitrate changes from MCS 0 to MCS 7.

As we explained in IV-B, ETX send probes only at the lowest STBC bitrate. As a result, in an 802.11n WMN, nodes see many neighbors connected over low-loss links at MCS 0; however, several of these links may perform poorly at higher bitrates. Choosing randomly among those neighbors results in shorter paths compared to an 802.11g WMN (Figure 6), consisting of fewer, longer, low-speed links. On the other hand, ETT takes link bandwidth into account in addition to loss rates, and hence, it should be able to distinguish high throughput SDM links from longer STBC links with similar loss rates. However, our results show that ETT selects similar paths to ETX. This suggests that the bandwidth estimation based on the Packet Pair technique is not accurate in an 802.11n WMN.

F. Link bandwidth

It is already known [11] that the accuracy of the popular Packet Pair technique suffers at the highest 802.11a/g bitrates

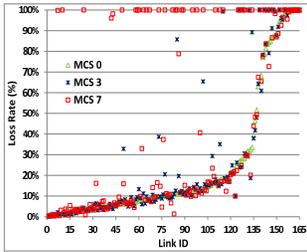
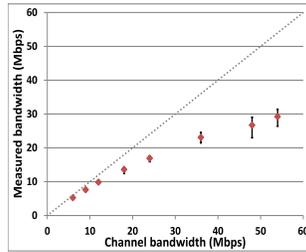
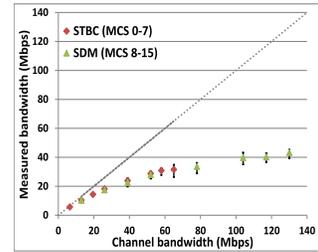


Fig. 10. Comparison of link loss rates for MCS 0, MCS 3, and MCS 7 in an 802.11n WMN.



(a) 802.11g.



(b) 802.11n.

Fig. 11. Accuracy of bandwidth estimation over 802.11g and 802.11n.

(36-54 Mbps). It is expected that the accuracy will drop even further at the highest 802.11n bitrates enabled by multiple MIMO streams and channel bonding. To verify this hypothesis, we measured the bandwidth of an 802.11g and an 802.11n link at all available bitrates using the Packet Pair technique. Two of our testbed nodes were placed near each other in the same room. We estimated the bandwidth of the wireless link between them using packet-pair probes.

The result is shown in Figures 11(a), 11(b) for 802.11g and 802.11n, respectively. Figure 11(a) confirms that the Packet Pair technique is quite accurate in the case of an 802.11g link and can unambiguously distinguish among different channel bandwidths, except, perhaps, for the two highest 802.11g bitrates. However, the accuracy drops significantly in the case of an 802.11n link (Figure 11(b)). The measured bandwidth is never higher than 45 Mbps, even though the channel bandwidth can be as high as 130 Mbps, and the Packet Pair technique can no longer unambiguously distinguish among various channel bandwidths. As a result, ETT ends up selecting randomly among links with similar loss rates at MCS 0, finally choosing paths similar to those chosen by ETX.

V. DISCUSSION

In this section, we discuss challenges in the design of link quality-based routing metrics in 802.11n WMNs and make a few recommendations for the next generation of routing metrics.

Loss rate estimation On one hand, probing all available bitrates [1] is not a viable option due to the increasing number of bitrates (16 for 2x2 MIMO, 24 for 3x3, etc). On the other hand, probing only the lowest bitrate and using the same estimate at higher bitrates [10], [11] may significantly reduce the accuracy (Figure 10). Hence, heuristics are needed for limiting probing to only a small subset of bitrates. The challenge is to identify for a given link the set of bitrates which are most likely to be used by the MAC RA algorithm. In fact, 802.11n RA algorithms face themselves a similar challenge [24], [14], and thus, they can provide useful insights in the design of efficient routing layer heuristics. One approach could be to use the signal-to-noise ratio (SNR) to identify a subset of the optimal bitrates and measure loss rates at only those bitrates [14]. A second approach which can be used alone or in combination with the first one, is adaptive probing, e.g., as done in [24], [26].

Link bandwidth estimation Figure 11(b) shows that the

traditional Packet Pair technique cannot provide accurate results for 802.11n links. We believe that a more promising solution is to approximate throughputs using the MAC bitrates, similar to in [1]. However, unlike in 802.11a/b/g networks, approximating throughput simply as the product of the bitrate and the delivery rate at that bitrate will not yield satisfactory accuracy for high bitrates in 802.11n. Two factors affect throughput and they have to be taken into account: MAC/PHY overhead and FA/BA which tries to amortize this overhead. MAC/PHY overhead is fixed per transmission and can be easily incorporated in throughput calculations. On the other hand, taking FA into account is not trivial, as the length of an A-MPDU depends on factors such as traffic type and link conditions, which vary over time.

FA/BA The throughput of an 802.11n link is much lower than the MAC bitrate without FA [27] due to the fixed MAC/PHY overheads involved in a frame transmission. The Packet Pair method cannot avoid these overheads and can severely underestimate link bandwidth, especially at high bitrates. On the other hand, approximating throughput simply as the product of the bitrate and the delivery rate at that bitrate implicitly assumes that FA/BA is always used and estimates link bandwidth to be close to the PHY bitrate.

In practice, there is no guarantee that FA will always be used. The number of aggregated frames into one A-MPDU exhibits a complex dependence on the link loss rate [24]. Moreover, for certain types of traffic (VoIP, Instant Messaging, HTTP), FA will simply not take place because node queues will never be backlogged. The consequence is that a metric like ETT in [1] may wrongly prefer lossy links operating at high bitrates, if it makes the assumption that lossless throughput over those links is always close to the bitrate. We believe that routing metrics for 802.11n WMNs should track the FA level (e.g., the average A-MPDU size) online and take it into account in the bandwidth estimation. We plan to further study this as part of our future work.

VI. RELATED WORK

Several link quality-based routing metrics have been proposed over the past decade both for single-channel [10], [17], [1], [28] and multi-channel networks [11], [29]. All these metrics have been designed assuming a legacy 802.11a/b/g underlying MAC layer and have only been evaluated over 802.11a/b/g WMNs.

The only metric that takes into account some of the new

characteristics introduced by the 802.11n standard (MAC/PHY overhead, FA) is the ECOT metric [28], which is calculated as the expected time occupancy during which a data frame or a group of data frames is transmitted over the expected number of successfully transmitted data frames at a unit transmission attempt. However, ECOT was only evaluated in ns2 over an 802.11g PHY layer.

Some works have also compared the performance of different routing metrics [16], [12], [30] in 802.11a/b/g WMNs. The only work that evaluated HC and ETX over a MIMO network is [31]. The main finding is that the gain of MIMO over SISO is 20% lower with ETX than with HC. The evaluation was done in simulations and only focused on the diversity provided by STBC. Our result in Figure 3 is qualitatively similar; 802.11n offers large throughput improvements over 802.11g with HC, small to mediocre improvements with ETX, and minimal or negative improvements with ETT.

Several researchers have experimentally studied the performance of the various features of 802.11n over the past few years [27], [32], [15], [33], [34], [35], [36]. All these studies focused on single-hop WLANs. A few recent works [37], [38], [39] examined multihop performance in 802.11n WMNs, studying the impact of the 802.11n features on UDP throughput and video streaming. A single chain with nodes close to each other was used for all the experiments. To our best knowledge, no work has investigated routing protocol performance in an 802.11n network.

VII. CONCLUSION

We conducted the first experimental study of link quality-based routing metrics in a 21-node indoor 802.11n WMN testbed. Our two main findings are (i) the observed throughput improvements of link quality-based routing metrics over the HC metric in 802.11g WMNs do not carry over in 802.11n WMNs and (ii) the large per-link throughput gains of 802.11n over 802.11a/b/g do not translate into similar multihop throughput gains. We further found that popular probing techniques for link loss rate and link bandwidth estimation yield poor accuracy in 802.11 WMNs due to the new features introduced at the underlying MAC/PHY layers. Finally, we observed that these metrics show a tendency towards short paths, consisting of long links, in 802.11n WMNs – a behavior similar to that of the HC metric in 802.11a/b/g WMNs. Our study shows the need for new routing metrics in WMNs over high-speed 802.11 radios. Our findings are not limited within the scope of 802.11n but are applicable to future 802.11 standards, such as 802.11ac [40], offering even higher bitrates and larger sets of MAC/PHY states.

ACKNOWLEDGMENT

We thank Naveen Rawat, Li Sun, and the anonymous reviewers for their helpful comments.

REFERENCES

- [1] J. Bicket, et al., “Architecture and evaluation of an unplanned 802.11b mesh network,” in *Proc. of ACM MobiCom*, 2005.
- [2] “Technology For All (TFA).” <http://tfa.rice.edu>.
- [3] “Meraki mesh routing.” <http://www.meraki.com/technologies/mesh-routing>.
- [4] Tropos systems, http://www.tropos.com/products/performance_mesh.html.
- [5] ARUBA networks, “Improving coverage and quality in video surveillance applications.” White Paper, 2011, <http://www.arubanetworks.com/technology/white-papers/#Mesh>.
- [6] J. Eriksson, et al., “A feasibility study of mesh networks for an all-wireless office.” in *Proc. of ACM MobiSys*, 2006.
- [7] “Mesh Networks for Home Automation.” <http://compnetworking.about.com/od/homeautomation/a/mesh-networking-in-home-automation.htm>.
- [8] D. Aguayo, et al., “Link-level measurements from an 802.11b mesh network.” in *Proc. of ACM SIGCOMM*, 2004.
- [9] B. Raman, K. Chebrolu, D. Gokhale, and S. Sen, “On the feasibility of the link abstraction in wireless mesh networks,” *IEEE/ACM Transactions on Networking (TON)*, vol. 17, no. 2, April 2009.
- [10] D. S. J. D. Couto, et al., “A high-throughput path metric for multi-hop wireless routing.” in *Proc. of ACM MobiCom*, 2003.
- [11] R. Draves, J. Padhye, and B. Zill, “Routing in multi-radio, multi-hop wireless mesh networks.” in *Proc. of ACM MobiCom*, 2004.
- [12] Y. Yang, J. Wang, and R. Kravets, “Designing routing metrics for mesh networks.” in *Proc. of WiMesh*, 2005.
- [13] “IEEE Std. 802.11n-2009: Enhancements for higher throughput.” <http://www.ieee802.org>, 2009.
- [14] K. LaCurts and H. Balakrishnan, “Measurement and analysis of real-world 802.11 mesh networks.” in *Proc. of ACM SIGCOMM/USENIX IMC*, 2010.
- [15] K. Pelechris, et al., “Experimental characterization of 802.11n link quality at high rates.” in *Proc. of ACM WINTECH*, 2010.
- [16] R. Draves, J. Padhye, and B. Zill, “Comparison of routing metrics for static multi-hop wireless networks.” in *Proc. of ACM SIGCOMM*, 2004.
- [17] S. Keshav, “A control-theoretic approach to flow control.” in *Proc. of ACM SIGCOMM*, 1993.
- [18] “UBMesh: An 802.11a/b/g/n wireless mesh network testbed at UB.” <http://www.cse.buffalo.edu/faculty/dimitrio/research/ubmesh/index.html>.
- [19] T. Clausen, et al., “Optimized link state routing protocol (OLSR).” RFC 3626, Oct 2003. [Online]. Available: <http://tools.ietf.org/html/rfc3626>
- [20] “OLSR routing protocol.” <http://www.olsr.org>.
- [21] “OLSR with Link Cost Extensions.” <http://sourceforge.net/projects/olsr-1c/>.
- [22] “iperf.” <http://dast.nlanr.net/Projects/Iperf>.
- [23] K. Sundaresan and R. Sivakumar, “Routing in ad-hoc networks with MIMO links.” in *Proc. of IEEE ICNP*, 2005.
- [24] I. Pefkianakis, et al., “MIMO rate adaptation in 802.11n wireless networks.” in *Proc. of ACM MobiCom*, 2010.
- [25] D. Nguyen and J. Garcia-Luna-Aceves, “A practical approach to rate adaptation for multi-antenna systems,” in *Proc. of IEEE ICNP*, 2011.
- [26] B. Radunovic, et al., “Dynamic channel, rate selection and scheduling for white spaces.” in *Proc. of ACM CoNEXT*, 2011.
- [27] V. Shrivastava, et al., “802.11n under the microscope.” in *Proc. of ACM/USENIX IMC*, 2010.
- [28] S. Kim, et al., “MAC-Aware routing metric for 802.11 wireless mesh networks.” in *Proc. of IEEE PIMRC*, 2009.
- [29] A. P. Subramanian et al., “Interference aware routing in multi-radio wireless mesh networks.” in *Proc. of IEEE WiMesh*, 2006.
- [30] S. M. Das, et al., “Studying Wireless Routing Link Dynamics.” in *Proc. of ACM SIGCOMM/USENIX IMC*, 2007.
- [31] E. Gelal, et al., “On the impact of MIMO diversity on higher layer performance.” in *Proc. of IEEE ICDCS*, 2010.
- [32] K. Pelechris, et al., “Design and deployment considerations for high performance MIMO testbeds.” in *Proc. of ICST WICON*, 2008.
- [33] M. Arslan, et al., “Auto-configuration of 802.11n WLANs.” in *Proc. of ACM/SIGCOMM CoNEXT*, 2010.
- [34] L. Deek, et al., “The impact of channel bonding on 802.11n network management.” in *Proc. of ACM/SIGCOMM CoNEXT*, 2011.
- [35] A. Z. und Robert Sombrutzki, “Evaluation of a low-cost IEEE 802.11n MIMO testbed.” Humboldt-Universität zu Berlin, SAR-TR-2011-12, 2011.
- [36] S. Lakshmanan, et al., “Realizing high performance multi-radio 802.11n wireless networks.” in *Proc. of IEEE SECON*, 2011.
- [37] S. Frohn, et al., “Analyzing the effective throughput in multi-hop IEEE 802.11n networks.” in *Proc. of IEEE HotMESH*, 2010.
- [38] J. Friedrich, et al., “Understanding IEEE 802.11n multi-hop communication in wireless networks.” in *Proc. of WinMee*, 2011.
- [39] S. Gubner and C. Lindemann, “Evaluating the impact of frame aggregation on video-streaming over IEEE 802.11n multihop networks.” in *Proc. of IEEE HotMESH*, 2012.
- [40] “IEEE p802.11ac. Specification framework for tgac. IEEE 802.11-09/0992r21”, January 2011.