Quantifying the Overhead due to Routing Probes in Multi-Rate WMNs

Ioannis Broustis*, Konstantinos Pelechrinis*, Dimitris Syrivelis‡, Srikanth V. Krishnamurthy*, Leandros Tassiulas‡
*University of California, Riverside ‡University of Thessaly
{*broustis, kpele, krish*}@*cs.ucr.edu* {*jsyr, leandros*}@*inf.uth.gr*

***Abstract*—The selection of high-throughput routes is a key element towards improving the performance of wireless multihop networks. While several routing metrics have been proposed in the literature, it has been shown that link-quality aware metrics can provide significantly higher end-to-end throughput. To date, the online computation of such metrics requires the periodic transmission of probe packets at all available transmission rates. However, our link level measurement study on two different 802.11 testbeds demonstrates that: (a) multi-rate probe transmissions increase the number of collisions and enforce nodes to reside in the back-off state for prolonged time periods, and (b) the extent of performance degradation depends on the network density; a network-wide throughput reduction of the order of *400%* is possible. In addition, our measurements show that the impact of probing in terms of end-to-end performance can be devastating. In particular, the probing functionality can pose a significant degradation in the end-to-end throughput of a single flow, by at least *35%* and as high as *90%*, depending on the probing frequency and network density. Finally, we discuss different alternatives to multi-rate probing for the online computation of such metrics.**

***Index Terms*—Wireless Mesh Networks (WMNs), Probing, Packet Delivery Ratio (PDR), Link Quality Aware Routing, Measurements.**

I. INTRODUCTION

Discovering high-performance paths in multihop networks is challenging and has received a lot of attention during the last few years. Typically a routing protocol relies on a performance metric, in order for efficient paths to be discovered. As an example, AODV [1] and DSR [2] rely on the *hop count* metric; such protocols discover routes with the smallest number of intermediate hops. Recent studies have shown that shortest-path routing may not always be efficient; short routes may be comprised of a few, low-SINR (Signal to Interference plus Noise Ratio) links, and this can degrade performance due to the poor link quality [3] and the resulting use of low modulation data rates [4], [5], [6], [7], [8]. As shown in [3], the PDR is a very accurate reflector of the link quality. Hence, modern routing protocols consider the PDR for estimating the quality of the individual links. In particular, such protocols try to find routes that consist of high-quality links [4], [5], [9]. In order to accomplish this, they compute the PDR of individual links using *probe* packets. The PDR values are then used to estimate the end-to-end quality of a route.

Probing is overhead intensive: Multi-rate probing, however, incurs a significant overhead. First, in order to accurately estimate the PDR, the size of a probe should correspond to the average data packet size (order of 1500 bytes). Second, since probes are broadcast packets, they are not supposed to be acknowledged; their transmission is usually prioritized over data packets [10]. These implementation decisions exacerbate the following two effects in 802.11-based wireless mesh networks:

1) *Packet collisions:* probe packets transmitted by a node often collide with data packets that are sent by neighbor nodes.
2) *Extensive back-off periods:* due to the CSMA/CA design principles, transmitters are coerced to reside in the back-off state for prolonged periods of time (since they classify the medium as busy more frequently).

Both these effects result in a reduction of the long-term throughput. Additionally, due to the dynamic nature of the wireless medium, the coherence time of the channel can be very small. Thus, the joint route and rate adaptation process has to be robust to frequent changes in link quality [11]. To address this requirement, the PDR should be measured in a timely way and with high accuracy, which implies that probe transmissions should be frequently invoked. However, this exacerbates the above two effects, as we discuss later.

As our main contribution in this paper, we perform measurements on different testbeds in order to quantify the amount of overhead that is imposed due to multi-rate probing. Our measurements involve both indoor and outdoor single-hop experiments, as well as experiments with multiple hops and routes of different lengths. Our study shows that the probing functionality that is currently integral to the routing operations can severely degrade the link throughput (e.g. by 75% as compared to single-rate probing with a 50 msec probe frequency) and the end-to-end throughput (e.g. by as much as 95%, depending on the frequency of probe transmissions as well as the network density). These observations suggest that a new method for computing the PDR is essential, which will not incur such high overheads. Towards developing such a method, we discuss different alternatives to probing that could potentially require offline measurements or accurate simulation modeling.

The remainder of the paper is structured as follows. In Section II, we provide the relevant background and discuss previous work. In Section III, we present our measurements for quantifying the performance degradation due to multi-rate probing. Our assessment on the effects of multi-rate probing involves experiments with single links as well as routes of different length. In Section IV, we discuss potential methods that can be used as an alternative to probing in link quality estimation. Finally, our conclusions form Section V.

II. BACKGROUND AND PREVIOUS STUDIES

In this section, we provide a brief background on the previously proposed routing metrics and discuss related work.

A. Link quality aware routing metrics

It has been previously shown that shortest path routing [1], [2] for wireless networks is not always efficient [4], [5], [6], [7], [8]. In particular, De Couto *et al.* [4] perform an extensive set of experiments, showing that the utilization of shortest path metrics in wireless networks may result in routes with significantly lower throughput than other potential routes with more hops. This is because shorter paths tend to have longer links, which are likely to operate in a low SINR regime. This affects the achievable throughput on the individual links on the route: a path with significantly lower end-to-end throughput than other potential routes with more hops is utilized.

De Couto *et al.* in [4] propose the *expected transmission count* routing metric, **ETX**. For a single link, ETX corresponds to the expected number of transmissions (including potential retransmissions) until the successful reception of the packet. For a route, the ETX is the sum over all the ETX values of the individuals links that comprise the route. The source node computes the ETX for all the potential routes to the destination, and selects the route with the minimum (aggregate) ETX value. The calculation of the ETX for a single link requires the online measurement of the forward (p_f) and the reverse (p_r) packet success probabilities on the link:

$$ETX = 1/(p_f \cdot p_r) \tag{1}$$

A variation of the ETX metric, called **ETOP** is proposed in [12]. The design of this metric is based on the consideration that the number of MAC layer retransmissions is not infinite. ETOP inherits the PDR calculation requirement of ETX.

Draves *et al* [5] propose the expected transmission time (**ETT**) metric, which is an adjustment of ETX. The ETT metric for a given link is the expected time to successfully send an MTU-size packet at that link's highest-throughput bit-rate, again including the time for any potential retransmissions. With ETT the source calculates the expected transmission time of each packet, taking into account the bandwidth of each link. ETT is calculated as:

$$ETT = ETX \cdot S/B \ , \tag{2}$$

where S is the packet size and B is the bandwidth of the link.

Finally, in 802.11s mesh networks, RM-AODV is the default routing protocol [13]. RM-AODV employs a new metric for routing, called **airtime cost.** This metric reflects the amount of channel resources that are consumed when transmitting a frame; it is calculated as:

$$C = [O_{ca} + O_p + B/R]/p. \tag{3}$$

In the above equation, O_{ca} is the channel access overhead, O_p is the protocol overhead, B is the number of bits in the probe packet, R is the current transmission rate and p is the PDR.

All these routing metrics rely on the online measurement of the PDR for each link. In order to be able to perform such a measurement, a *probing functionality* is incorporated into the implementation of each of these metrics. We discuss the basics of the probing functionality below.

B. Sending probes to measure the PDR

To date, the probing functionality in link quality aware routing typically operates as follows. Each node in the network transmits broadcast packets periodically every τ seconds. These packets are called probes, and their size usually reflects the size of the packets that travel over the air. For example, the size of probes that correspond to data packets are usually of the same order as the MTU size. Since ACK is disabled for broadcast packets, each node reports the number of probe packets received by each of its neighbors, during a *probing window* of w seconds. This information is piggybacked on to future probes, and provides information to other nodes with regards to the forward and the reverse packet success probabilities *(ratio of probes that were actually received over the number of probes expected within* w*)* for different links. Specifically with ETT, each node sends periodic 1500-byte broadcasts (reflecting data packets) *at each available 802.11 bit-rate*, and periodic 60-byte broadcasts (reflecting ACK frames) at the basic rate. As discussed above, this information is used by each node to compute the values of p_f and p_r and further the value of ETT as per equation 2. Clearly, the frequent transmission of these probes incurs additional overhead, which can impact the network operations. The goal of our work is to quantify these overheads and their impact on network performance. We also discuss alternatives to the probing functionality.

There are some recent studies [3], [14], [15] that aim in evaluating the currently available ways of estimating link quality. They exhibit interesting findings, however they are orthogonal to our measurement study.

III. QUANTIFYING THE OVERHEADS DUE TO MULTI-RATE PROBING

In this section, we discuss our experimental observations with regards to the performance degradation that is induced due to probing. Throughout the rest of the paper we consider ETT-based routing; our observations will be analogous with any other metric that is based on the online computation of the PDR.

A. Motivating our measurements

As discussed in Section II, with ETT each node is periodically transmitting probes at different bit rates in order

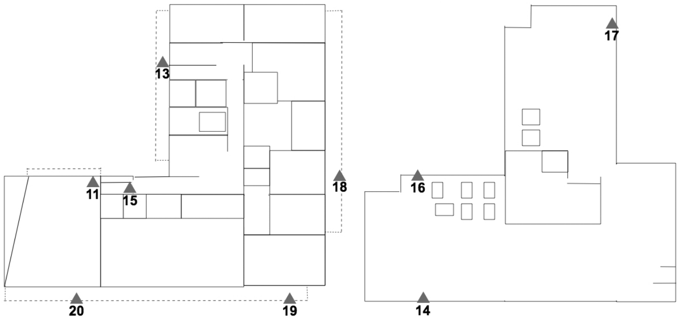

Fig. 1. Our 802.11a/g testbed, deployed on the 4th floor (left) and the rooftop (right) of the ECE building, at the University of Thessaly.

to compute the PDR with each bit rate. However, there is a limitation on how fast the channel information can be obtained. The authors in [9] have set the frequency of probe transmissions (probing interval) to once every 1 sec, while the packet error probability is calculated every 10 sec. Note however that the *coherence time* of a channel can be quite small [11], i.e., the channel conditions may vary at very short time scales. Therefore, it is highly likely that the computation of the PDR through probing is not accurate when executed at low-frequency intervals, such as on the order of seconds or hundreds of milliseconds. This suggests that probes must be sent out much more frequently at every rate, and that the PDR should be calculated potentially every few milliseconds. However, this can affect the throughput at single links as well as the overall network performance. In what follows, we present a set of throughput measurements that support our arguments.

B. Effects of probing on link throughput

We first discuss our experiments with single-hop scenarios, where we observe the effects of probing on the performance of individual links.

Experimental network description: We conduct this set of experiments on our 9-node wireless testbed, which we have deployed in 2 floors of the ECE building at the University of Thessaly. The node layout is depicted in Fig. 1. The testbed has been mainly deployed outdoors and only 2 nodes (namely 11 and 15) are placed indoors. There is connectivity between the two floors, i.e., some nodes from the 4th floor can communicate with nodes on the rooftop. The nodes are based on the ORBIT hardware configuration [16], and run a Debian Linux distribution with kernel v2.6.16.19 over NFS. Each node is equipped with 1 GHz CPU, 512 Mbytes of memory, and a WN-CM9 wireless mini-PCI card, which carries the AR5213 Atheros main chip. We utilize the RoofNet toolkit [17], which hosts the full implementation of the ETT metric and the SRCR routing protocol [9]. We use the MadWifi v0.93 driver; we have modified the driver to cooperate with the Roofnet software.

Measurement study: We examine how the probing functionality at different bit rates on a single link impacts the throughput on individual neighbor links. For this, we differentiate the links under examination into (a) "victim" links, which are ***not*** employing any form of probing, and (b) "probing" links; we call the nodes that send probes as "probers". In most of our experiments, victim links are set to a fixed transmission rate; we experiment with different rates. We perform our experiments with 802.11g, late at night, in order to avoid interference from collocated WLANs. We consider the following cases, while we experiment with different traffic scenarios, topologies and link qualities.

a. One victim link and one prober: We enable fully-saturated UDP traffic with 1500-byte data packets on a single victim-link. Meanwhile we activate an interfering (to the victim) prober. There is no data traffic transmitted by the latter; the prober is simply sending probes at all available rates (8 rates for 802.11g). With this experimental configuration we are able to observe the actual impact of multi-rate probing on the performance of single victim links.

We demonstrate the effects of multi-rate probing for the scenario where the victim link is 15→11 and the prober is node 20 (locations shown in Fig. 1). The 15→11 link is a high-SINR link, which achieves 27.9 Mbits/sec when operating in isolation. In Fig. 2, we observe that the impact on performance from a single prober is quite significant. While the accurate determination of the packet error probability requires frequent probing, this reduces the throughput on neighbor victim links to a large extent. Moreover, the throughput reduction is more severe as we increase the probing frequency. On the other hand, the transmission of probes at a single rate[1] affects the performance of the victim link to a much smaller extent, as shown in Fig. 2. This is because (a) the victim link observes that the medium is idle more frequently thereby managing to transmit more packets per unit time, and (b) the number of collisions is now reduced. While this is somewhat expected (only 1/8^{th} of the probes are now sent out) it clearly provides an intuition about the performance benefits that are incurred by a more flexible packet error rate calculation scheme that would not require multi-rate probing.

Furthermore, we conduct experiments with victim links that perform rate adaptation; the transmission rate per link is selected to be the rate with the lowest ETT. We observe that probing affects the performance even with rate adaptation. Our measurements indicate that the ETT implementation typically tends to utilize high transmission rates, such as the 54 Mbits/sec rate. While this rate selection provides high long-term throughput, it is accompanied by low PDR values. In the Appendix we provide a simple analytical explanation of why the ETT design may many times prefer such high bit rates.

b. One victim link and many probers: We further seek to observe how the introduction of more probing links affects the performance of a victim link. For this, we increasingly activate more probers in the neighborhood of 15→11, one by one. We observe that:

---

[1] We have used 54 Mbps as the single probing rate. This corresponds to the lowest possible probing overhead, since the flight time of the probe at 54 Mbps is the shortest as compared to the time with the other rates.

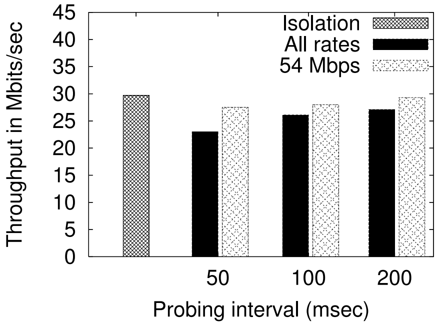

Fig. 2. The frequency of probing invocation affects the performance of a single victim link.

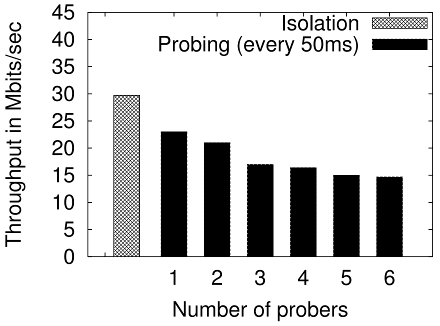

Fig. 3. Increasing the number of prober nodes progressively degrades the performance.

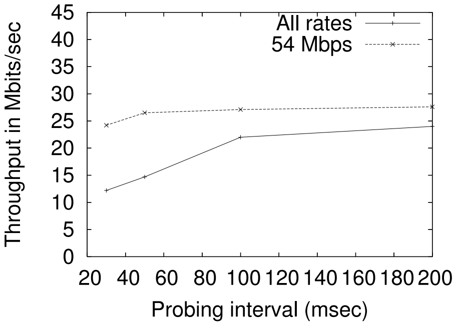

Fig. 4. Throughput on link 15→11 with 6 interfering probers and different probing schemes.

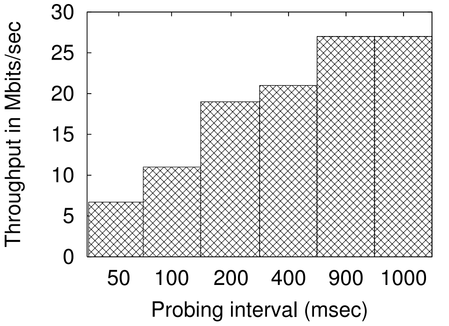

Fig. 5. Multi-rate probing leads to very poor performance with higher densities.

1) As we increase the number of probers, the throughput at the victim link is further reduced. As an example in Fig. 3, for 5 (multi-rate) interfering probers we observe a throughput reduction by 50.5%.
2) Similarly as before, the frequency of probing plays an important role in the achieved performance: the higher the frequency, the lower the throughput, as we observe in Fig. 4.
3) Single-rate probing does not seem to affect the performance as much. As depicted in Fig. 4, when probes are sent at 54 Mbits/sec only, the throughput on 15→11 is reduced by 15% only, with the 50-msec probing frequency.

c. Case where all links use probing: We also experiment with scenarios where all nodes are sending both traffic and probes. For this, we consider 4 links that form a multi-hop topology: links 11→20, 13→15, 16→14 and 18→17. We enable the ETT functionality on all nodes, and we initiate single-hop saturated UDP traffic on each link. Our measurements are depicted in Fig. 5, where the total network throughput (sum of throughput of individual links) is plotted. From Fig. 5 we notice that our aforementioned observations also hold in scenarios where all nodes are transmitting probes and data. Since we perform our experiments late at night in a totally controlled environment, the ETT mechanism is likely to find appropriate transmission rates with much less frequent probing intervals. As we observe in Fig. 5, when probes are transmitted once per 900 ms at every rate, the total network throughput is 27 Mbits/sec. This throughput is equal to the one achieved when probes are transmitted once per 1000 msec; this is also equal to the probe-free throughput. *We observe that multi-rate probing every 50ms can degrade the total throughput by **400%***! Clearly one may expect an even more severe performance reduction in higher density topologies and/or with shorter probing intervals. This experimental study demonstrates that while frequent multi-rate probing can be used to accurately compute the packet loss probability, it introduces tremendous delay overheads in the network, thereby degrading performance.

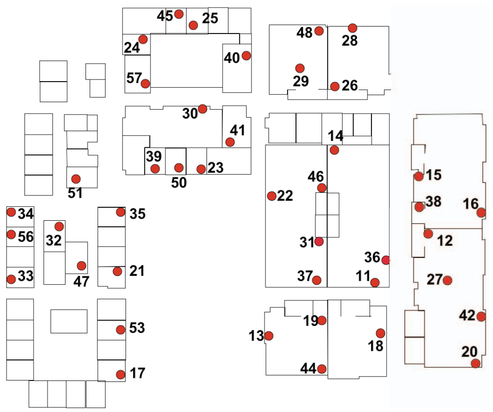

Fig. 6. The deployment of our indoor testbed at the University of California, Riverside. Nodes are represented by dots along with their IDs.

C. Effects of probing on the end-to-end throughput

Next, we perform measurements to assess the impact of multi-rate probing on the end-to-end performance.

Experimental network description: We conduct this set of experiments on a different, large-scale testbed. This testbed has the exact same software configuration as the previous one, however it is deployed in the 3^{rd} floor of the Engineering Building Unit II, at the University of California, Riverside. This testbed differs from our testbed at the Univ. of Thessaly in the following aspects:

(a) It is much larger: it consists of 42 nodes, and this enables us experiment with more routing scenarios. **(b)** It is all deployed indoors (see Fig. 6), and the environmental conditions differ substantially from the deployment in Fig. 1.

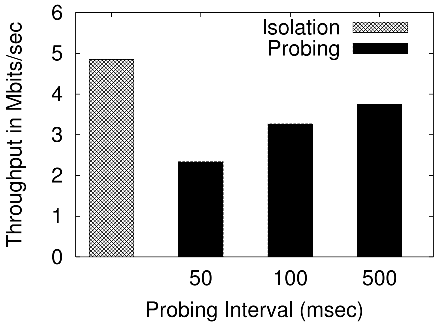

Fig. 7. Frequent probing can degrade the performance by up to 50% even with a single prober.

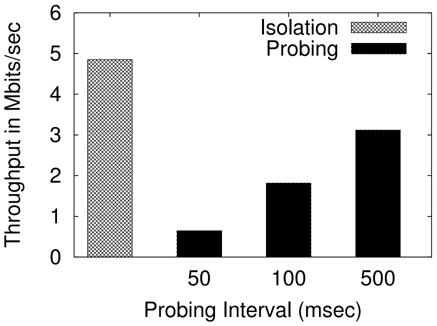

Fig. 8. Increasing the number of probers can further degrade the performance (3 probing nodes).

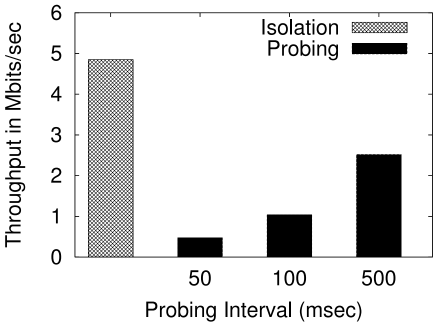

Fig. 9. Higher densities result in more profound degradation due to probing (5 probing nodes).

(c) The hardware in this testbed is also different. In particular, the testbed consists of 22 Soekris net5501 boxes and 20 Soekris net4826 boxes [18], however they carry the same, Roofnet-compatible WiFi card [17].

Measurement study: We opt to examine the effects of multi-rate probing on the end-to-end performance. We use a large set of high-throughput routes, hardcoded on our testbed using the `route` Linux tool, which manipulates the kernel routing table on each node. For this set of experiments we use SampleRate [19] as the rate adaptation algorithm. Here we discuss our observations for routes consisting of 2, 3 and 4 hops.

We initiate traffic on the flows in isolation (no other flow is active at the same time) and we measure the end-to-end throughput, for different populations of probing nodes in the vicinity of the flow and for different probing intervals. We experiment with each route in isolation, in order to capture the actual effects of probing on performance, avoiding the interference that would be incurred by other active routes. Fig. 7, 8 and 9 present our results for the 3 hop routes on our testbed that were able to achieve approximately 5 Mbps end-to-end throughput in isolation. We observe that the presence of a single, frequent prober (probing interval at 50 msec), can significantly degrade the performance, i.e., by as much as 50% (throughput drops from 4.85 Mbps in isolation, to just 2.34 Mbps with probing). Recall that the nodes participating in the traffic session do not use probes (they use fixed routing tables), and that the probers are simply transmitting probes, without participating in forwarding data traffic. The performance degradation is more profound if we increase the number of probers. In particular, for the scenario of 5 probers, even for the case of infrequent probing (probing interval 500 msec), the throughput degradation is approximately 50% (achievable throughput 2.52Mbps). With a higher probing frequency (i.e. probing interval 50 msec) the throughput degradation on the route becomes more severe (approximately 90%). Our results with different configurations (number of hops and achievable end-to-end throughput in isolation) exhibit the same qualitative performance.

Our experimental results indicate that the presence of multi-rate probing can significantly degrade the end-to-end performance, in addition to the individual link performance. This artifact suggests that a new method to estimate the PDR on a link is required. We discuss the possible operation of such mechanisms in the following section.

IV. DISCUSSING ALTERNATIVES TO MULTI-RATE PROBING

In this section, we discuss alternative techniques to multi-rate probing. Here we simply discuss the potential operation of those mechanisms; we intend to implement and evaluate them in our future work.

The design of such alternative PDR estimation techniques could involve the calculation of the PDR based on its relationship to the SNR (Signal-to-Noise-Ratio) for the different transmission rates. More specifically, at any given instance the SNR at the receiver is the same, irrespective of the bit rate that is used by the transmitter. In other words, the use of different bit rates does not affect the reception SNR. However, the PDR is typically different at each transmission rate: the higher the rate, the lower the PDR [20]. In addition, note that at each specific rate there is a PDR value that corresponds to every SNR value [20]. Hence, *if the reception SNR is known, the PDR can be estimated for any available transmission rate*. This artifact can be exploited by PDR estimation strategies that do not rely on multi-rate probing.

Creating simulation-based models: One such strategy could be to develop a simulation model, which would provide a large set of PDR-SNR tuples (e.g. simulated PDR points at different SNR values, perhaps in 1dB or 2 dB steps) for each rate. Such a simulation PHY model would potentially involve random bit generation, convolutional encoding, bit-to-QAM symbol mapping, Rayleigh fading generation, additive white Gaussian noise (AWGN) generation, symbol decision, log-likelihood ratio (LLR) computation and Viterbi decoding. Such a simulation model could provide a dense set of PDR values for different SNR regimes and different transmission rates. Hence, if the PDR at a given rate is known, this strategy could utilize the model to:

1) retrieve the corresponding SNR (provided by the PHY simulations), and
2) estimate the PDR for every other available bit rate, since the SNR is the same for all those rates.

In order to compute the PDR at a certain rate (and use it to perform 1 and 2 above), the transmission of periodic single-rate probes could suffice. In other words, with such a mechanism the multi-rate probing could be replaced by single-rate probing, thereby reducing the incurred overhead dramatically. Clearly, the estimated PDR as per this approach could deviate from the actual PDR, since this whole approach is based on a simulation model. On the other hand, the application of this mechanism could provide performance benefits due to the lower incurred probing overheads. We plan to investigate this accuracy vs. performance trade-off in our future work.

Use of offline measurements: A pertinent approach could rely on actual offline measurements in order to derive the aforementioned PDR-SNR tuples, instead of developing a PHY simulation model. Such an approach could be performed in two different ways. First, a large set of PDR values could be collected offline for different transmission rates, in order to construct the aforementioned PDR-SNR tuples. With this, the above steps (1 and 2) would be similarly performed to estimate the PDR at any available bit rate; the only difference with this approach is that the pre-stored PDR-SNR tuples will have been measured offline (instead of being derived through PHY layer simulations). Second, the reception SNR could be directly measured (as we explained earlier, the SNR is the same at any transmission rate). A similar strategy is proposed in [21], where Verma *et al.* perform offline measurements to derive PDR vs. SNR mappings. However, measuring the SNR with a high accuracy is hard and this can hurt the overall system performance.

V. CONCLUSIONS AND FUTURE WORK

Multi-rate probing is the most popularly considered approach for estimating link quality for discovering high-throughput end-to-end paths. Frequently transmitted probes can incur a large number of data packet collisions and force nodes reside at the backoff state for extended periods. The effects of multi-rate probing become more and more prominent, as the density of the network deployment increases. Our measurements on two different wireless testbeds demonstrate that the network performance in terms of link and end-to-end throughput is significantly affected due to the frequent transmission of probes. Given the above artifacts, we discuss potential alternatives to multi-rate probing. Such alternatives could simply require probing at one rate only, in order to estimate the PDR any other rate through offine measurements or accurate simulation modeling. In our future work, we plan to implement these strategies in our testbeds and evaluate their efficacy, by comparing them against the multi-rate probing technique.

REFERENCES

[1] E. M. Royer and C. E. Perkins. An implementation study of the AODV routing protocol. In *IEEE WCNC*, 2000.

[2] D. B. Johnson, D. A. Maltz, and J. Broch. DSR: The Dynamic Source Routing Protocol for Multi-Hop Wireless Ad Hoc Networks. In *Ad Hoc Networking, Ch. 5, pp. 139-172*, 2001.

[3] A. Vlavianos, E. Law, I. Broustis, S. V. Krishnamurthy, and M. Faloutsos. Assessing Link Quality in IEEE 802.11 Wireless Networks: Which is the Right Metric? In *IEEE PIMRC*, 2008.

[4] D. S. J. De Couto, D. Aguayo, J. Bicket, and R. Morris. A High Throughput Path Metric for MultiHop Wireless Routing. In *ACM MOBICOM*, 2003.

[5] R. Draves, J. Padhye, and B.Zill. Routing in Multi-Radio, Multi-Hop Wireless Mesh Networks. In *ACM MOBICOM*, 2004.

[6] R. Dube, C. Rais, K.-E. Wang, and S. Tripathi. Signal Stability Based Adaptive Routing (SSA) for Ad-hoc Mobile Networks. In *IEEE Personal Comm*, 1997.

[7] T. Go, N. Abu-Aahazaleh, D. Phatak, and R. Kahvecioglu. Preemptive Routing in Ad Hoc Networks. In *ACM MOBICOM*, 2001.

[8] Y.-C. Hu and D. B. Johnson. Design and Demonstration of Live Audio and Video Over Multi-hop Wireless Networks. In *MILCOM*, 2002.

[9] J. Bicket, S. Biswas D. Aguayo, and R. Morris. Architecture and Evaluation of an Unplanned 802.11b Mesh Network. In *ACM MOBICOM*, 2005.

[10] Click web page. http://read.cs.ucla.edu/click.

[11] J. Camp and E. W. Knightly. Modulation Rate Adaptation in Urban and Vehicular Environments: Cross-Layer Implementation and Experimental Evaluation. In *ACM MOBICOM*, 2008.

[12] G. Jakllari, S. Eidenbenz, N. Hengartner, S. V. Krisnamurthy, and M. Faloutsos. Link Positions Matter: A Noncommutative routing Metric for Wireless Mesh Networks. In *IEEE INFOCOM*, 2008.

[13] 802.11s: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Simple Efficient Extensible Mesh (SEE-Mesh) Proposal.

[14] S.M. Das, H. Pucha, K. Papagiannaki, and Y. C. Hu. Studying Wireless Routing Link Metric Dynamics. In *ACM IMC*, 2007.

[15] M. Carrera, H. Lundgren, T. Salonidis, and C. Diot. Correlating Wireless Link Cost Metrics to Capacity. In *IEEE/IFIP WONS*, 2009.

[16] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh. Overview of the ORBIT Radio Grid Testbed for Evaluation of Next-Generation Wireless Network Protocols . In *IEEE WCNC*, 2005.

[17] MIT Roofnet. http://pdos.csail.mit.edu/roofnet.

[18] Soekris Engineering. http://www.soekris.com.

[19] J. Bicket. Bit-rate Selection in Wireless Networks. In *MS Thesis, Dept. of Electr. Engin. and Comp. Science, MIT*, 2005.

[20] H. Jafarkhani. *Space-Time Coding, Theory and Practice*. Cambridge University Press, 2005.

[21] L. Verma, S. Kim, S. Choi, and S.J. Lee. Reliable, Low Overhead Link Quality Estimation for 802.11 Wireless Mesh Networks. In *IEEE WiMesh*, 2008.

APPENDIX

The ETT metric for a single link is calculated as per Eq. 2. ETT selects the bit rate that will provide the highest throughput in the long term. However, as we show here, this rate selection may provide poor performance in terms of the Packet Delivery Ratio (PDR).

Let us examine the two extreme cases of 6 and 54 Mbits/sec for a link; the corresponding ETT values are ETT_6 and ETT_{54}. From Eq. 1 and 2, we have $PDR = p_f \cdot p_r$ and $ETT = \frac{1}{PDR} \cdot \frac{S}{B}$. Recall that the rate with the lowest ETT value will be selected. We compare the ETT values for 6 and 54 Mbits/sec. In particular we examine in which cases ETT_6 will be smaller than ETT_{54}:

$$ETT_6 < ETT_{54} \;\Leftrightarrow\; \frac{1}{PDR_6} \cdot \frac{S}{6} < \frac{1}{PDR_{54}} \cdot \frac{S}{54} \;\Leftrightarrow$$

$$PDR_{54} < \frac{6}{54} \cdot PDR_6$$

From the above inequality, we observe that in order for the 54 Mbps rate **not** to be selected, the corresponding PDR must be extremely low. As an example, if we assume that PDR_6=0.8 and PDR_{54}=0.1, then the rate of 54Mbps is going to be selected, since the above inequality does not hold. In other words, the highest bit rate will be selected, although the PDR at this rate is extremely low.