On the Efficacy of Opportunistic Routing

Zifei Zhong* and Srihari Nelakuditi†
* University of Texas, Austin † University of South Carolina, Columbia

***Abstract*—Traditional routing schemes select the *best* path for each destination and forward a packet to the corresponding next hop. While such *best-path routing* schemes are considered well-suited for networks with *reliable point-to-point* links, they are not necessarily ideal for wireless networks with *lossy broadcast* links. Consequently, *opportunistic routing* schemes that exploit the broadcast nature of wireless transmissions and dynamically select a next-hop *per-packet* based on loss conditions at that instant are being actively explored. It is generally accepted that opportunistic routing performs substantially better than best-path routing for wireless mesh networks. In this paper, we analyze the efficacy of opportunistic routing. We define a new metric EAX that captures the expected number of any-path transmissions needed to successfully deliver a packet between two nodes under opportunistic routing. Based on EAX, we develop a candidate selection and prioritization method corresponding to an ideal opportunistic routing scheme. We then conduct an off-line comparison of best-path routing and opportunistic routing using our EAX metric and MIT Roofnet trace. We observe that while opportunistic routing offers better performance than best-path routing, the gain is not as high as commonly believed.**

I. INTRODUCTION

There have been a variety of routing protocols proposed for multi-hop wireless networks. Despite their many differences, a common aspect of most of these schemes is that they attempt to find the *best* path and forward packets to the corresponding next hop. When a packet is lost due to transmission errors, they either retransmit it to the same next hop or rediscover a new best path. Such *best-path routing* schemes are likely to trigger many packet retransmissions or path rediscoveries since wireless transmissions tend to have high loss rates as they are susceptible to external interference, multi-path fading and inclement weather. Moreover, wireless channel conditions vary at a fast time scale that the best path at an instant may not be good at the next instant. Therefore best-path routing which is considered well-suited for wired networks with relatively stable point-to-point links, may not be an ideal approach for wireless networks with lossy broadcast links.

Consequently, *opportunistic routing* schemes that mitigate the impact of lossy channels by exploiting the broadcast nature of wireless transmissions are being actively explored [1]–[3]. Extremely opportunistic routing (ExOR) is one such hop-by-hop routing scheme initially proposed in [1] and later developed into a source routing scheme in [2]. The focus of this paper is on hop-by-hop routing and so here we describe the operation of ExOR of [1] but later discuss the features of ExOR of [2] also. The general idea behind ExOR is that, for each destination, a set of next-hop candidates are selected and each of them is assigned a priority according to its *closeness* to the destination. When a packet needs to be forwarded, the highest priority node is chosen as the next-hop, *after* the packet's transmission, among the candidates that received it. Thus, in contrast to the best-path routing, where a packet is unicast to the predetermined next-hop, under opportunistic routing, a next-hop is determined per-packet after its broadcast transmission.

It is envisioned that opportunistic routing reduces the number of transmissions needed for reliable delivery of a packet, as it avoids retransmissions as long as the packet makes forward progress towards the destination. However, it runs the risk of duplicate forwarding by multiple candidates unaware of others' transmission resulting in potentially more overall number of transmissions than even best-path routing. Therefore, the utility of opportunistic routing hinges on inter-candidate communication in ensuring that only the highest priority candidate that received the packet forwards it. Towards this end, ExOR makes all the candidates relay the acknowledgements by the higher priority candidates effecting robust acknowledgements and thus avoiding both unnecessary retransmissions by the senders and duplicate forwarding by the candidates. It is reported that ExOR improves throughput by more than a factor of two over best-path routing particularly for distant node pairs with more than two hops [1], [2]. But it is not clear how much of this gain is exclusively due to the opportunistic selection of next-hops since other features such as robust acknowledgements and scheduling of transmissions could play a major part in the overall performance.

To assess the true benefit of opportunistic routing, in this paper, we compare it with best-path routing which also employs robust acknowledgements. We make several assumptions in our study. We suppose that there is only one flow in the network and the aim is to minimize the average number of transmissions needed to deliver a packet of that flow. We further imagine that there are no collisions and packet loss is only due to the channel conditions eliminating the role of scheduling of transmissions. We also assume that both best-path and opportunistic routing do not use RTS and CTS frames as in [1], [2]. We consider only those schemes that select a next-hop after DATA transmission such as ExOR unlike others such as [3] that select a next-hop after receiving a CTS but before sending the DATA frame. Though all these assumptions narrow the focus of our study, we believe it still reveals valuable insights into the utility of opportunistic routing for wireless mesh networks.

Our study employs the following methodology offline given a wireless network topology and the corresponding link-level data frame delivery probabilities. We approximate the robust acknowledgement delivery probabilities by assuming that each acknowledgment is piggybacked on a data frame repeatedly a certain number of times. We define a new metric *expected any-path transmissions* (EAX) that captures the expected number of transmissions needed to successfully deliver a packet between a node pair under opportunistic routing, given the link-level data and acknowledgment delivery probabilities, and the set of candidates and their priorities. For best-path routing which has just one candidate, EAX boils down to ETX [4]. Based on EAX, we select an ideal set of candidates and prioritize them yielding the smallest possible EAX, for each node pair, under any opportunistic routing scheme. We then compare the ideal EAX of opportunistic routing with that of best-path routing.

The contributions and findings of this paper are as follows. We present a new routing metric EAX and a candidate selection procedure based on EAX suitable for opportunistic routing. We show that our approach yields fewer candidates and transmissions than ETX-based approach. The key finding of this paper is that while opportunistic routing offers performance improvement over best-path routing, the gain is much less than generally assumed. Moreover, robust acknowledgements reduce transmissions more than opportunistic forwarding. We analyze the reasons behind less-than-expected performance of opportunistic routing and observe that there are not many equally good candidates and long jumping links particularly at high data rates in Roofnet. These observations indicate that the utility of opportunistic routing depends critically on the density of wireless network. As cautioned before, we note that these findings are based on above assumptions and one network. Therefore, further thorough investigation is needed to confirm the benefits and limits of opportunistic routing.

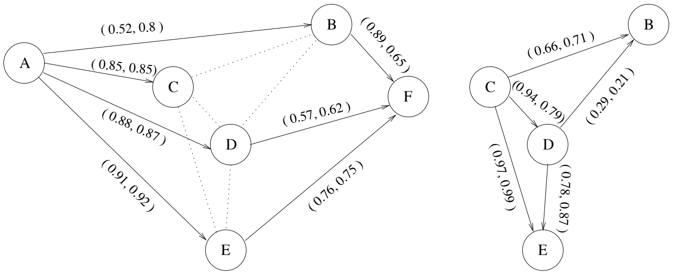

Fig. 1. A subgraph of Roofnet [7] used for illustrations. Each link is labelled with (f_1, f_2), where f_1 is the packet delivery probability in the direction of the arrow and f_2 is that in the other direction. For readability, the complete subgraph is shown on the left and the delivery probabilities between the intermediate nodes on the right.

The rest of this paper is organized as follows. Section II describes the methodology used in computing EAX for a node pair and selecting candidates based on EAX. Section III evaluates and compares the performance of ETX- and EAX-based opportunistic routing, and best-path routing. Section IV further analyzes the evaluation results. Section V discusses related work. Section VI concludes and discusses future work.

II. METHODOLOGY

There are several opportunistic routing schemes proposed for wireless networks such as [1], [2], [5], [6]. In this paper, we focus specifically on opportunistic hop-by-hop routing and describe in this section the methodology we used in evaluating its efficacy. We first introduce our model of opportunistic routing framework and illustrate how we compute the probability that a candidate nexthop forwards a packet it receives. We then show the computation of EAX metric for a node pair, given the set of candidates and their priorities for all node pairs. Finally, we demonstrate how candidates are chosen and assigned priorities based on EAX.

A. Opportunistic Routing Framework

First, we assume that the probability of delivering an ACK in a single transmission is the same as that of a data packet, which is the case when ACKs are piggybacked on data packets. We model *robust acknowledgment* (RACK) mechanism like ExOR's *batch* [2] as if an ACK is repeated a certain number of times, referred

TABLE I
ACK DELIVERY PROBABILITY WITH ROBUST ACK MECHANISM

RACK size	C↔E	C↔D	C↔B	D↔E	D↔B
1	(0.97, 0.99)	(0.94, 0.79)	(0.66, 0.71)	(0.78, 0.87)	(0.29, 0.21)
2	(1.00, 1.00)	(1.00, 0.96)	(0.88, 0.92)	(0.95, 0.98)	(0.50, 0.38)
5	(1.00, 1.00)	(1.00, 1.00)	(1.00, 1.00)	(1.00, 1.00)	(0.82, 0.69)
10	(1.00, 1.00)	(1.00, 1.00)	(1.00, 1.00)	(1.00, 1.00)	(0.97, 0.91)
20	(1.00, 1.00)	(1.00, 1.00)	(1.00, 1.00)	(1.00, 1.00)	(1.00, 0.99)

to as the RACK size. It is also assumed that the inclusion of RACK would not increase the length of a data packet significantly. Therefore, if the delivery probability of a packet is a, then the *RACK delivery probability* with RACK size n is given by $1 - (1 - a)^n$. For example, consider the topology shown in Figure 1, where each link is labelled with the corresponding packet delivery probabilities in each direction. Suppose A is the source, F is the destination and the other four nodes are the candidates for forwarding. The delivery probabilities between candidates is shown separately on the right. The RACK delivery probabilities corresponding to the links between candidates is shown in Table I for different RACK sizes. Hereafter, whenever we refer to the delivery probability of an ACK, we mean RACK delivery probability.

We model opportunistic packet forwarding process as follows. Each node selects a subset of its neighbors as candidate next-hops for a destination and assigns a priority to each of the candidates. When a sender transmits a packet, each candidate that received the packet responds, in turn according to its priority, with a RACK indicating the highest priority candidate, known to this candidate, that received the packet. Note that only those candidates that received the packet are involved in acknowledging and forwarding it. A sender retransmits a packet only if it does not receive a RACK from any of the candidates. A candidate forwards the packet if it is the highest priority candidate that it knows to have received the packet. Duplicate forwarding by two candidates is possible if a lower priority candidate can not hear a RACK either directly or indirectly from a higher priority candidate. The probability of that happening is captured by the following formulation.

Let s be the source and d be the destination. Suppose $C^{s,d}$ is the set of candidate next-hops from s to d, and c_i is the candidate with priority i (with 1 being the highest). Assume that the packet delivery probability from s to c_i is f_i and RACK delivery probability from c_i to s is a_i. Similarly, let a^i_j be the probability of RACK delivery directly from c_j to c_i. A candidate c_i can get informed of a higher priority candidate c_j's reception of the same packet either from c_j directly or indirectly through another lower priority candidate c_k $(k > i)$ that also received the packet from s and the RACK from c_j. When the set of candidates is large, there could be many levels of indirection such as c_j to c_{k_1} to c_{k_2} to c_{k_3} to c_i through which c_i gets to know that c_j also received the packet. If we limit the level of indirection to just one, we can approximate λ_i, the probability that the i-th priority candidate does not get informed of any higher priority candidates' reception of the same packet, as follows.

$$\lambda_i = \prod_{j<i} \left(1 - f_j + f_j(1 - a^i_j) \prod_{k>i} (1 - a^k_j a^i_k f_k) \right) \quad (1)$$

Here, f_j is the probability that a higher priority candidate c_j receives the packet, and $f_j(1 - a^i_j)$ is the probability that c_i does not get informed of it by c_j. Similarly, $(1 - a^k_j a^i_k f_k)$ is the probability that c_i is not informed of c_j's reception by a lower priority candidate c_k that received the packet. Therefore, the probability that c_i does not get informed of c_j's reception either by c_j or any of the lower priority candidates is $f_j(1 - a^i_j) \prod_{k>i}(1 - a^k_j a^i_k f_k)$. Hence, $1 - f_j + f_j(1 - a^i_j) \prod_{k>i}(1 - a^k_j a^i_k f_k)$, is the probability of c_j not receiving the packet or c_i not receiving c_j's acknowledgement. Thus, equation (1) captures the probability that c_i does not get informed of any higher priority candidates' reception of the packet.

The above formula gives the exact value of λ_i when the size of the candidate set is less than four, otherwise it does not account for the propagation of RACK to i from a candidate with higher priority than i indirectly through two or more candidates that have priorities lower than i. However, even with four or more candidates, it yields a reasonably good approximation of λ_i. For illustration, Table II lists the approximate (as per the above equation) and exact values of λ corresponding to the candidates B,E,C,D (ordered according to their priority highest to lowest) in Figure 1. It can be observed that equation (1) yields the exact λ values for candidates B, C, and D. Only when calculating it for E, equation (1) does not take into account the possibility that a RACK is propagated from B to E through first from B to C, then C to D and finally D to E. However, this only makes approximated λ value slightly higher than its actual value. The difference between approximate and actual becomes negligible as the RACK size increases, because in that case with very high probability a candidate gets a RACK from at least one higher priority candidate. When the RACK size is large enough to have 100% RACK delivery probability, duplicate forwarding can be completely eliminated, in which case, $\lambda_i = \prod_{j<i}(1 - f_j)$.

TABLE II

λ: ACTUAL VS. APPROXIMATION

RACK size	B (λ_1)		E (λ_2)		C (λ_3)		D (λ_4)	
1	1.00	1.00	**0.648**	**0.665**	0.057	0.057	0.037	0.037
2	1.00	1.00	**0.558**	**0.559**	0.046	0.046	0.013	0.013
5	1.00	1.00	0.511	0.511	0.043	0.043	0.009	0.009
10	1.00	1.00	0.496	0.496	0.043	0.043	0.007	0.007
20	1.00	1.00	0.490	0.490	0.043	0.043	0.007	0.007

TABLE III

DISTANCE TO F ACCORDING TO THE ETX OF THE BEST-PATH AND THE EAX WITH OPPORTUNISTIC ROUTING.

RACK	A		B		C		D		E	
size	ETX	EAX	ETX	EAX	ETX	EAX	ETX	EAX	ETX	EAX
1	2.95	2.90	1.73	1.73	2.80	2.79	2.83	2.13	1.75	1.75
2	2.51	2.45	1.28	1.28	2.45	2.43	2.05	1.70	1.40	1.40
5	2.42	2.27	1.13	1.13	2.35	2.33	1.77	1.59	1.32	1.32
10	2.42	2.25	1.12	1.12	2.35	2.32	1.75	1.59	1.32	1.32
20	2.42	2.24	1.12	1.12	2.35	2.25	1.75	1.59	1.32	1.32

B. Expected Any-path Transmissions

We now define the expected number of any-path transmissions needed for reliable delivery of a packet from a source s to a destination d, given the candidate set $C^{s,d}$, as follows,

$$\text{EAX}(s,d) = \mathcal{S}(s,d) + \mathcal{Z}(s,d) \tag{2}$$

$$\mathcal{S}(s,d) = \frac{1}{1 - \prod_i (1 - f_i a_i)} \tag{3}$$

$$\mathcal{Z}(s,d) = \frac{\sum_i \lambda_i f_i \text{EAX}(c_i, d)}{1 - \prod_i (1 - f_i)} \tag{4}$$

where $\mathcal{S}(s,d)$ captures the expected number of transmissions for successfully transmitting a packet from s to at least one of the candidate next-hops and getting at least one acknowledgment back to s, and $\mathcal{Z}(s,d)$ captures the expected number of transmissions for delivering the packet in turn from those candidates to the destination. When the RACK size is large leading to reliable RACK delivery, it can be expressed as

$$\text{EAX}(s,d) = \frac{1 + \sum_i \text{EAX}(c_i, d) f_i \prod_{j=1}^{i-1}(1 - f_j)}{1 - \prod_i (1 - f_i)} \tag{5}$$

The difference between the two metrics ETX and EAX in estimating the distance from each node to destination F of Fig. 1 is shown in Table III. The ETX values shown correspond to the best path to F from each node. In case of EAX, we use the procedure described in the next section for selecting candidates and list the corresponding EAX values. The difference between ETX and EAX values indicates the extent of gain possible with opportunistic routing over best-path routing.

C. Candidate Selection and Prioritization

It is possible that the candidate selection based on an inappropriate metric can actually degrade the performance of opportunistic routing. For example, under ExOR two next-hop nodes c_1 and c_2 are both selected as candidates by source s if the ETX distances from c_1 and c_2 to destination d are both closer than that from s to d. This could lead to duplicate forwarding by c_1 and c_2 in case these two candidates can not communicate with each other at all. The proposed new metric EAX accounts for the potential duplicate forwarding by the candidates and helps determine the *contribution of a candidate* to the delivery of packets between a node pair and thus enables judicious selection of candidates.

Alg 1 : **SELECT**(s, d)

```
 1: {Lines 2-5: generate candidate pool $\mathcal{G}$}
 2: $\mathcal{G} \Leftarrow \emptyset$
 3: for all $v_i \in \mathcal{N}(s)$ do
 4:     if $\text{ETX}(v_i, d) < \text{ETX}(s, d)$ then
 5:         $\mathcal{G} \Leftarrow \mathcal{G} \cup \{v_i\}$
 6: {Lines 7-22: pick contributing candidates $\mathcal{C}$ from $\mathcal{G}$}
 7: $m_p \Leftarrow \infty$
 8: $m_c \Leftarrow \infty$
 9: $\mathcal{C} \Leftarrow \emptyset$
10: while TRUE do
11:     {Lines 12-15: find the next best candidate $v$}
12:     for all $v_j \in \mathcal{G}$ do
13:         if $m_c > \text{EAX}(\mathcal{C} \cup \{v_j\}, s, d)$ then
14:             $v \Leftarrow v_j$
15:             $m_c \Leftarrow \text{EAX}(\mathcal{C} \cup \{v_j\}, s, d)$
16:     {Lines 17-20: update the candidate list $\mathcal{C}$}
17:     if $m_c < m_p$ then
18:         $\mathcal{C} \Leftarrow \mathcal{C} \cup \{v\}$
19:         $\mathcal{G} \Leftarrow \mathcal{G} \setminus \{v\}$
20:         $m_p \Leftarrow m_c$
21:     else
22:         break {no more qualified candidates}
23: return $\mathcal{C}$
```

The `SELECT` procedure for selecting candidates based on EAX at a node s for a specific destination d is shown as Alg 1. It is important to note that EAX between a node pair depends on the set of candidates, whereas their selection in turn is based on EAX. Therefore, the candidate selection is an iterative refinement process and this SELECT procedure is repeated with all node pairs till there is no change in the candidate set of any node pair. Let $\mathcal{N}(s)$ be the set of neighbor nodes of s, $\mathcal{G}$ the set of potential candidates, and $\mathcal{C}$ the set of actual candidates. Assume that $\text{ETX}(s,d)$ returns the ETX distance from s to d, and $\text{EAX}(\mathcal{C}, s, d)$ returns the

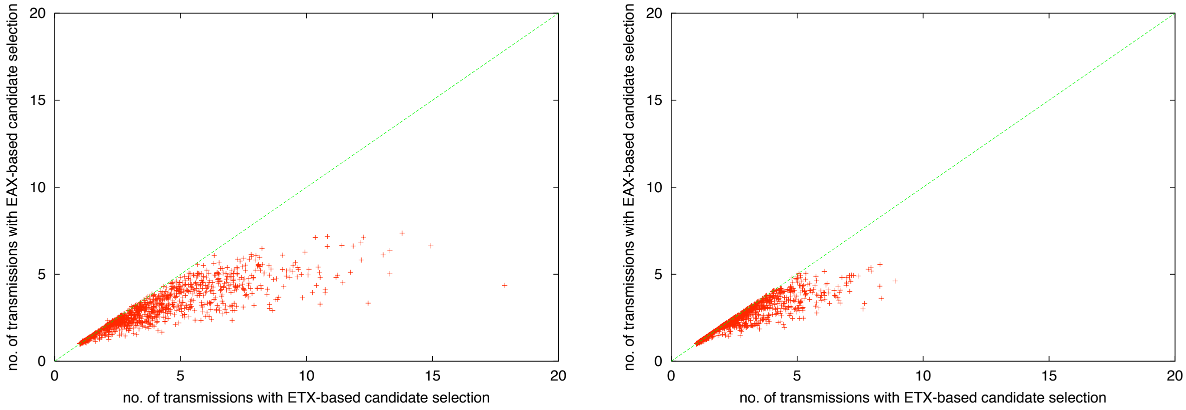

(a) 1 Mbps, RACK size 1 (b) 1 Mbps, RACK size 10

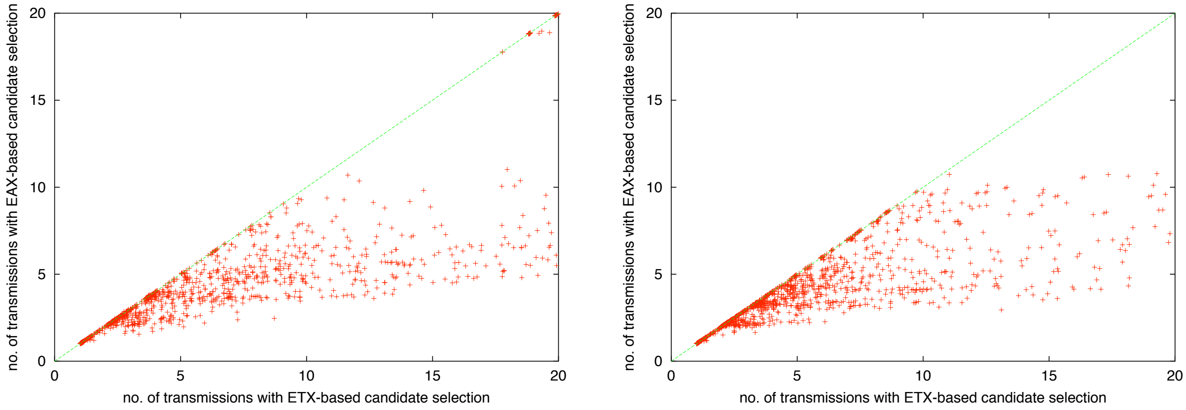

(c) 11 Mbps, RACK size 1 (d) 11 Mbps, RACK size 10

Fig. 2. Comparison of performance of opportunistic routing with candidate selection based on ETX and EAX. In all the cases, EAX-based candidate selection results in a fewer expected number of transmissions for a packet delivery between any node pair.

EAX distance from s to d with the candidate set $\mathcal{C}$. In `SELECT`, lines 2-5 generate the candidate pool $\mathcal{G}$, which is determined based on the best path ETX. A neighbor v_i is included in $\mathcal{G}$ only if ETX(v_i, d) $<$ ETX(s, d). The candidate set $\mathcal{C}$ is initialized to empty and then a subset of $\mathcal{G}$ is incrementally added to $\mathcal{C}$. Note that the candidate selection for all the nodes in $\mathcal{G}$ is done before it is done for s. Lines 12-15 find the best candidate in $\mathcal{G}$ that reduces EAX(s, d) the most and it is added to the set $\mathcal{C}$ in lines 17-20. This process in lines 10-20 is repeated till no new candidates are added to $\mathcal{C}$. Once the set of candidates are determined, they are prioritized according their EAX to the destination, i.e., the candidate with the smallest EAX is assigned the highest priority.

The differences in candidate selection based on ETX and EAX is illustrated using Fig. 1. Assume that node A is the source, node F is the destination, and the RACK size is 1. Based on ETX, node A will choose 4 candidates: B, E, C, and D, because paths from these nodes to F have smaller ETX than that from A to F. With EAX-based approach, only 3 candidates, E, D, and C, will be selected by source A, since adding B to the candidate set does not decrease EAX between A and F. This difference in candidate selection leads to different number of transmissions: 3.24 with ETX-based selection and 2.90 with EAX-based selection. The set of candidates and priorities for different RACK sizes are summarized in Table IV. Note that even when the candidate sets are same for A, with RACK sizes beyond 5, the number of transmissions are different under ETX- and EAX-based selection because the set of candidates of B, E, C, and D in turn (not shown) are different. Note that in this example, opportunistic routing with ETX-based selection of candidates results in more transmissions than even single best-path routing which needs only 2.95 transmissions (refer Table III). This illustrates the potential harm caused by opportunistic routing in the form of duplicate forwarding due to imperfect coordination between the candidates, and thus emphasizes the importance of appropriate candidate selection.

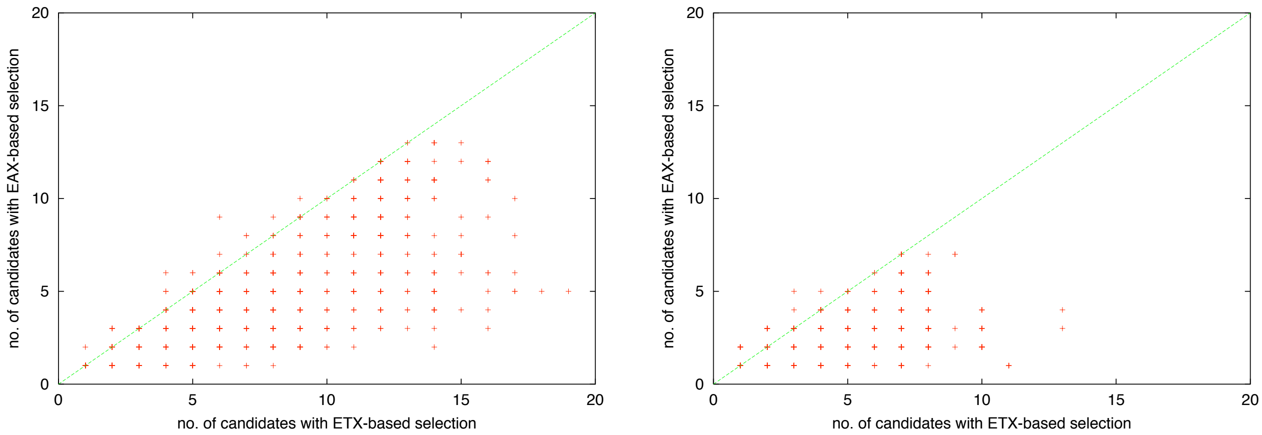

(a) 1 Mbps (b) 11 Mbps

Fig. 3. Comparison of the number of candidates selected based on ETX and EAX with RACK size 10. These plots show that these approaches select different set of candidates and ETX-based approach is much more likely to select more candidates.

TABLE IV
ETX- VS. EAX-BASED CANDIDATE SELECTION

RACK size	metric	candidates	# Tx
1	ETX	B>E>C>D	3.24
	EAX	E>D>C	2.90
2	ETX	B>E>D>C	3.19
	EAX	E>D>C	2.45
5	ETX	B>E>D>C	3.19
	EAX	B>E>D>C	2.27
10	ETX	B>E>D>C	3.19
	EAX	B>E>D>C	2.25
20	ETX	B>E>D>C	3.19
	EAX	B>E>D>C	2.24

III. EVALUATION

We now evaluate the efficacy of opportunistic routing using the link-level measurement trace of MIT Roofnet [7], a 38-node multi-hop wireless mesh network. The measurement trace records packet delivery over each link of the network for a total of 90 seconds with transmitting rates: 1, 2, 5.5, and 11 Mbps. We compute the average delivery ratio over 90 seconds for each link, and use this average value as its link-level delivery probability. The resulting topology is used to compare ETX-based candidate selection with EAX-based candidate selection for opportunistic routing, and opportunistic routing with best-path routing.

A. ETX- vs EAX-based Candidate Selection

The process of selecting candidates and assigning priorities is a key component of any opportunistic routing scheme. We argued in the previous section that candidate selection and prioritization based on ETX is not appropriate as it does not account for opportunistic forwarding towards the destination in turn by the candidates, and duplicate forwarding due to imperfect RACK delivery between the candidates. We now show that the proposed EAX-based candidate selection takes into account both these aspects and yields better performance.

Fig. 2 compares the performance of opportunistic routing with ETX-based candidate selection and EAX-based candidate selection. It presents a scatter plot with the expected number of transmissions for delivering a packet between each node pair under ETX-based approach on x-axis and that under EAX-based approach on y-axis. The plots are shown for two rates 1 and 11 Mbps and two RACK sizes 1 and 10. It is evident that in all the cases, EAX-based candidate selection results in a fewer number of transmissions. The number of candidates chosen by these two approaches is shown in Fig. 3. It is clear that they select a different set of candidates and for most node pairs, ETX-based approach selects more candidates than EAX-based approach. These results confirm that EAX-based candidate selection is ideal for opportunistic routing as it reduces the number of transmissions using fewer candidates for forwarding.

B. Opportunistic Routing vs Best-path Routing

We now compare the performance of best-path routing (BR) and opportunistic routing (OR) with EAX-based candidate selection. To make the comparison fair, it is assumed that best-path routing also employs robust acknowledgments as does opportunistic routing. Fig. 4 contrasts the performance of best-path routing and opportunistic routing using scatter plots with the expected number of transmissions for delivering a packet between each node pair under best-path routing on x-axis and that

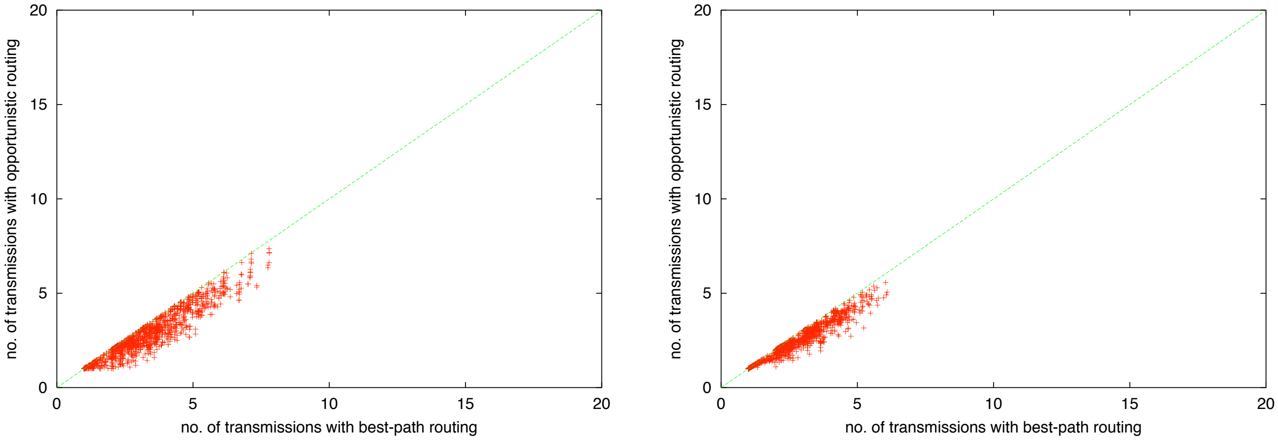

(a) 1 Mbps, RACK size 1

(b) 1 Mbps, RACK size 10

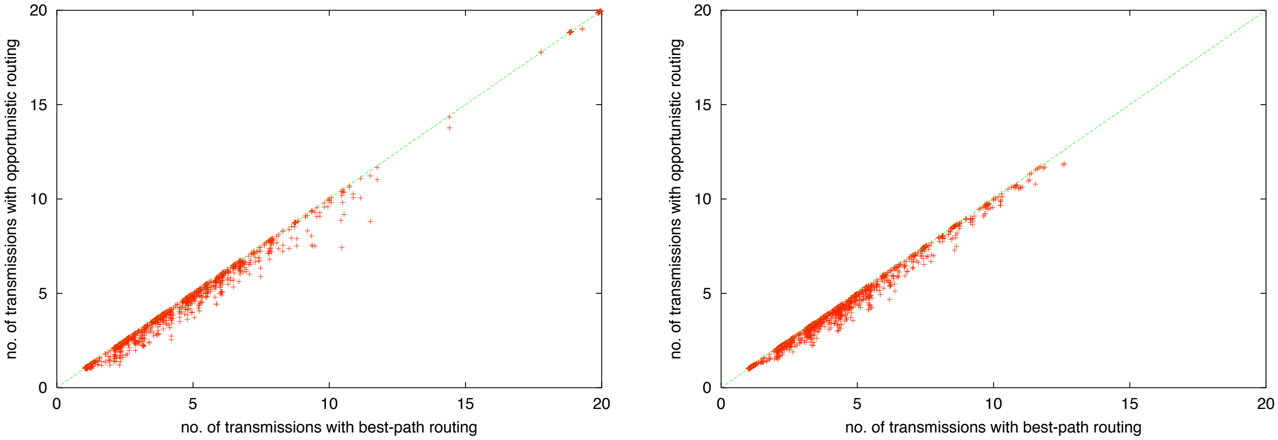

(c) 11 Mbps, RACK size 1

(d) 11 Mbps, RACK size 10

Fig. 4. Comparison of best-path routing and opportunistic routing. Opportunistic routing is always better but the gain is not quite high.

under opportunistic routing on y-axis. The plots are again shown for two rates 1 and 11 Mbps and two RACK sizes 1 and 10. It can be seen that in all the cases, opportunistic routing delivers with fewer transmissions than best-path routing but the gain is not as high as generally expected.

To investigate the effect of robust acknowledgements on the performance of these approaches, we plot the average number of transmissions per node pair under different RACK sizes in Fig. 5. The performance of the best-path routing with RACK size 1, i.e., the traditional routing without robust ACKs is shown as the horizontal line for reference. The difference in performance between BR with RACK size 1 and BR with RACK size 20 reflects the gain solely due to robust ACKs. On the other hand, the difference between BR and OR, both with RACK size 20, indicates the gain due to opportunistic forwarding. It is surprising to note that the gain due to robust ACKs is more than the gain due to opportunistic forwarding. This is even more pronounced at a higher data rate of 11 Mbps. This shows that comparing opportunistic routing that employs robust ACKs against traditional best-path routing without robust ACKs could unfairly inflate the utility of opportunistic routing.

It is expected that the node pairs that are several hops away have more to gain from opportunistic routing than those that are closer. To explore this, we relate the number of hops between a node pair and the corresponding performance gain under opportunistic routing in Fig. 6. It shows the relative average improvement of opportunistic routing w.r.t. best-path routing as a function of the number of hops along the best-path. There are several things to note here. Without robust ACKs, the relative performance of opportunistic routing does not seem to depend on the number of hops. On the other hand, with robust ACKs, it has a better relative performance in general when the number of hops are high. However, at 11 Mbps data rate, regardless of the setting, opportunistic routing offers no more than 5% improvement.

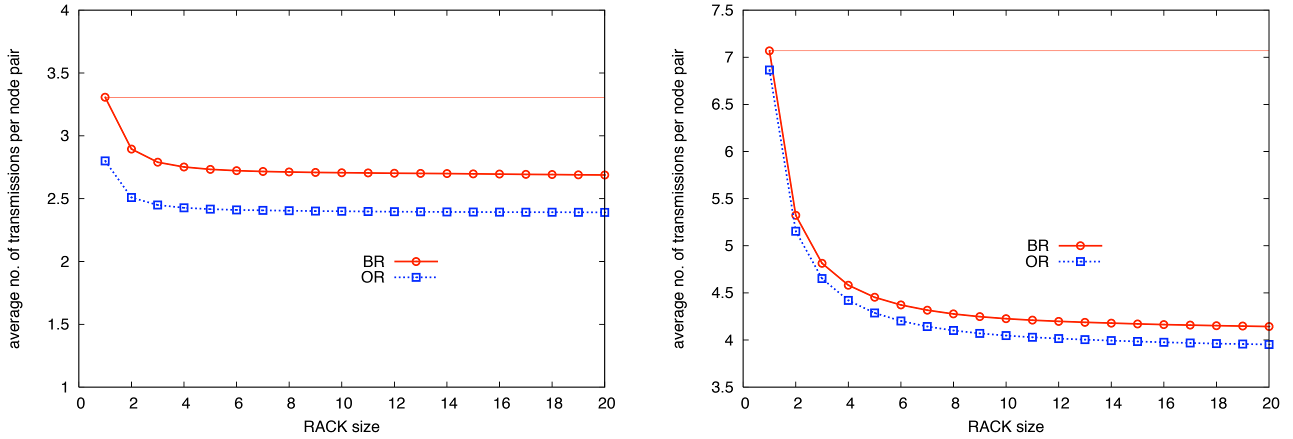

(a) 1 Mbps (b) 11 Mbps

Fig. 5. RACK size vs. average number of transmissions per node pair under best-path routing and opportunistic routing. Gain due to robust ACKs is more than the gain due to opportunistic forwarding.

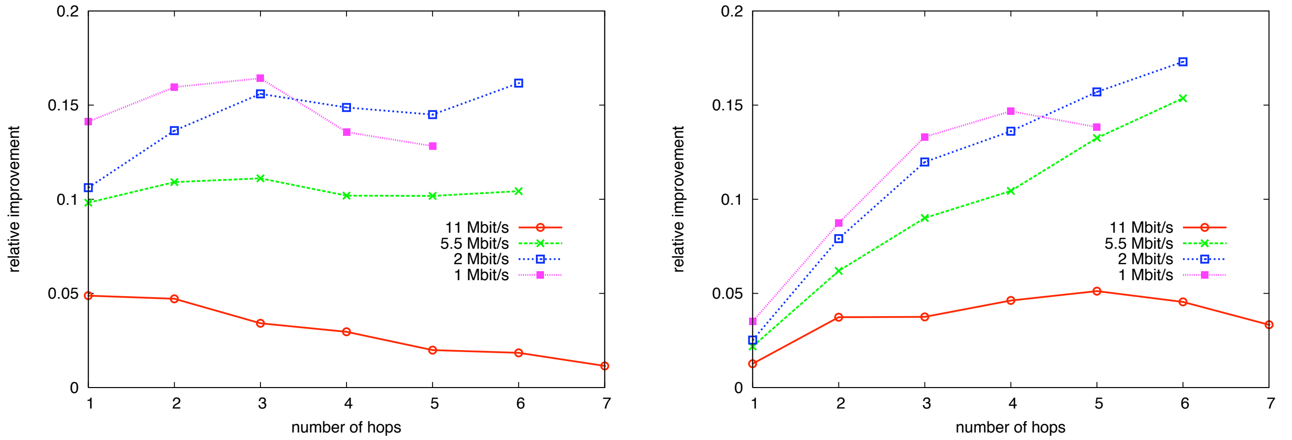

(a) RACK size 1 (b) RACK size 10

Fig. 6. Hop count vs. the relative improvement of opportunistic routing over best-path routing.

IV. ANALYSIS

We now analyze the causes for the performance of opportunistic routing to be less than expected for Roofnet. It is generally believed that opportunistic routing can exploit the existence of many good candidates making forward progress and many imperfect links skipping several hops. It is also argued that opportunistic routing is better when the loss probability is high and thus more suitable at high data rates. In this section, we investigate whether they hold true in case of Roofnet.

One of the illustrations used to demonstrate the potential of opportunistic routing is the one where there exist many equally good candidates effectively reducing the number of transmissions to 2 from 101 [2]. We define a candidate B as good if it satisfies the condition ETX(s,A)+EAX(A,d) - EAX(B,d) $>$ 0.5 * ETX(s,A), where A is the next-hop along the best path. In other words, a candidate is considered good if it makes forward progress at least half that of the best next-hop. We counted the number of such candidates for each node pair under each data rate and shown them in Fig. 7. These results show that majority of the node pairs of Roofnet do not have more than one good candidate.

Another scenario where opportunistic routing is touted to be good is when there exist many links that are not reliable enough for traditional best-path routing but are suitable for opportunistic forwarding as they make long jumps towards the destination. To investigate this, we first filter out those links with very low delivery probability, i.e., links with ETX of more than 100 (with delivery probability less than 10% in each direction). We then count a link $i-j$ as an opportunistic link, if the best-

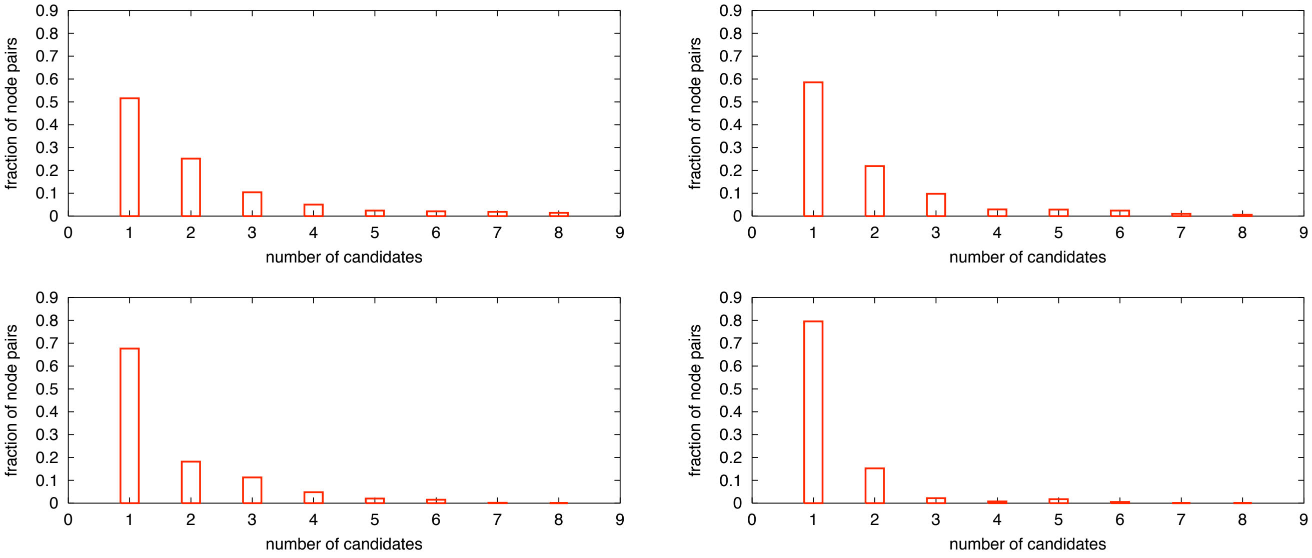

Fig. 7. This figure show the distribution of fraction of node pairs that have many equally good candidates. Y axis show the fraction, X axis shows the number of candidates. Each bar shows the fraction of node pairs that have that many equally good candidates.

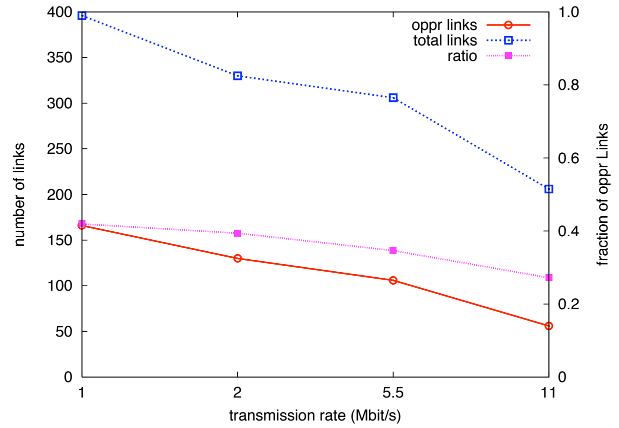

Fig. 8. The number of opportunistic links out of all the links under each transmitting rate. The y-axis on left shows the number of links and the y-axis at right shows the fraction of opportunistic links.

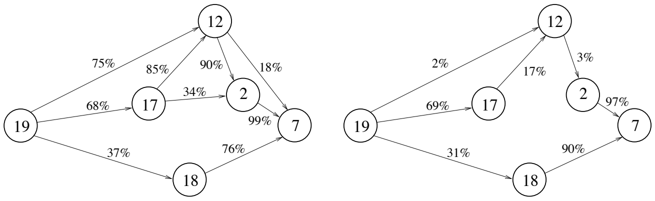

(a) 1 Mbps (b) 11 Mbps

Fig. 9. Illustrations extracted from Roofnet trace

path from i to j is not $i{-}j$. Fig. 8 shows the number of total links and opportunistic links under each rate. It seems that less than 40% of the links are opportunistic which may not be enough to offer sufficient opportunities for making long jumps towards the destination.

It is generally believed that opportunistic routing performs better at high rates of transmission since higher rate incurs higher loss and thus more potential for opportunistic forwarding. However, as we have seen in the previous section, opportunistic routing may not necessarily perform well at high transmission rates. This is due to the reduction in the number of links in the network as shown in Fig. 8. To illustrate this, we extracted two subgraphs from the Roofnet which are shown in Fig. 9. It shows that the topology corresponding to 11 Mbps has two fewer links and many links with very low delivery probabilities. Due to low density of the network, there will not be enough good candidates for opportunistic forwarding to perform well at high transmission rates.

V. RELATED WORK

There have been several opportunistic routing and forwarding schemes proposed for wireless networks which are discussed briefly below. This paper differs from existing work as it focuses on investigating the efficacy of opportunistic routing for wireless mesh networks.

OAR [8] attempts to identify the best channel and let that pair of nodes send without interruption, while maintaining fairness between contending senders. OAR works when channels are stable enough that many packets can be sent between changes in channel conditions. SDF [5] makes each sender indicate potential forwarders in the RTS packet. In SDF, candidate forwarders may send

CTS simultaneously and cause collisions. [9] improves the SDF approach by having the forwarders respond according to a priority order specified in the RTS frame. In MAC-layer Anycast [3] protocol, the knowledge of instantaneous channel condition is utilized to select next-hop relay on short time scales. It eliminates control packets, and forwarding decision is made based on historical observations of channel at MAC layer. Packet salvaging approach [10] uses opportunistic forwarding at MAC layer to enable packet salvaging whenever collision happens. Packet salvaging allows better channel spatial reuse by having a higher carrier sense threshold. SOAR [11] is designed to handle multiple flows in a wireless mesh network. ROMER [12] balances between long-term route stability and short-term opportunistic performance. OPRAH [13] is designed specifically for dynamic ad hoc networks that is robust to changes in network environment and mobility of ad hoc nodes.

GeRaF [6] is an opportunistic geographic routing protocol that tries to select the best positioned nodes as relays via contention among receivers. Each node that receives a message can assess its own priority based on location information and may volunteer to act as relay according to its priority. A framework for modeling such geographic region-based opportunistic routing protocols is presented in [14]. It assumes that there is no cross traffic at any node and separates out the routing, sleep discipline and medium access components to allow easy analysis. In [15], the authors investigate the factors that could affect the performance of opportunistic routing. However, their focus is on power consumption of geographic opportunistic routing schemes in sensor networks. The focus of this paper is on ExOR-like schemes that attempt to minimize the number of transmissions for reliable delivery in wireless mesh networks.

MORE [16] utilizes sophisticated network coding to avoid forwarding duplication without requiring tight coordination between candidates. Their work is motivated by the similar arguments as ours and we believe that MORE approach can also benefit from the EAX-based candidate selection presented in this paper.

VI. CONCLUSIONS

In this paper, we conducted an off-line comparison of best-path routing and opportunistic routing using our EAX metric and MIT Roofnet trace. We observed that opportunistic routing based on single path metric such as ETX could degrade performance. Instead a metric such as EAX that accounts for inter-candidate communication is more suitable for opportunistic routing. Moreover, we found that the benefit of opportunistic routing is much less than commonly believed. In general, the contribution of robust acknowledgements in reducing the number of transmissions seem to be more than that of opportunistic forwarding, especially at high transmission rates. Since these findings are based on a single network with several assumptions, we intend to investigate this further using simulations and real testbeds, and identify the scenarios where opportunistic routing is most effective.

REFERENCES

[1] S. Biswas and R. Morris, "Opportunistic Routing in Multi-hop Wireless Networks," in *HotNets-II*, Nov. 2003.

[2] ——, "Opportunistic Routing in Multi-hop Wireless Networks," in *ACM SIGCOMM*, Aug. 2005.

[3] R. R. Choudhury and N. H. Vaidya, "MAC-layer Anycasting in Ad Hoc Networks," *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 1, pp. 75–80, 2004.

[4] D. D. Couto, D. Aguayo, B. A. Chambers, and R. Morris, "Performance of Multihop Wireless Networks: Shortest Path is Not Enough," in *HotNets-I*, Oct. 2002.

[5] P. Larsson, "Selection Diversity Forwarding in a Multi-Hop Packet Radio Network with Fading Channel and Capture," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 5, no. 4, pp. 47–54, 2001.

[6] M. Zorzi and R. Rao, "Geographic Random Forwarding (GeRaF) for Ad Hoc and Sensor Networks: Multi-Hop Performance," *IEEE Transaction on Mobile Computing*, vol. 2, no. 4, Oct. 2003.

[7] MIT Roofnet, "http://www.pdos.lcs.mit.edu/roofnet/." [Online]. Available: http://www.pdos.lcs.mit.edu/roofnet/

[8] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic Media Access for Multirate Ad Hoc Networks," in *Proc. ACM MobiCom*, Sep. 2002.

[9] S. Jain and S. Das, "Exploiting Path Diversity in the Link Layer in Wireless Ad Hoc Networks," in *the 6th IEEE WoWMoM Symposium*, Jun. 2005.

[10] C. Yu, K. G. Shin, and L. Song, "Link-Layer Salvaging for Making Routing Progress in Mobile Ad Hoc Networks," in *Proc. ACM MobiHoc*, May 2005.

[11] E. Rozner, J. Seshadri, Y. Mehta, and L. Qiu, "Simple Opportunistic Routing Protocol for Wireless Mesh Networks," in *WiMesh*, 2006.

[12] Y. Yuan, H. Yang, S. H. Wong, S. Lu, and W. Arbaugh, "Romer: Resilient opportunistic mesh routing for wireless mesh networks," in *The 1st IEEE Workshop on Wireless Mesh Networks (WiMesh)*, Sep. 2005.

[13] C. Westphal, "Opportunistic Routing in Dynamic Ad Hoc Networks: the OPRAH protocol," in *IEEE MASS*, 2006.

[14] R. C. Shah, S. Wietholter, A. Wolisz, and J. M. Rabaey, "Modeling and Analysis of Opportunistic Routing in Low Traffice Scenarios," in *The 3rd International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, Apr. 2005.

[15] ——, "When Does Opportunistic Routing Make Sense?" in *The 1st International Workshop on Sensor Networks and System for Pervasive Computing*, Mar. 2005.

[16] S. Chachulski, M. Jennings, S. Katti, and D. Katabi, "A network coding approach to opportunistic routing," MIT, Tech. Rep. MIT-CSAIL-TR-2006-049, Jun. 2006.