Opportunistic Routing in Multi-Hop Wireless Networks

Sanjit Biswas and Robert Morris
MIT Laboratory for Computer Science
{biswas,rtm}@lcs.mit.edu

Abstract

This paper describes Extremely Opportunistic Routing (ExOR), a new unicast routing technique for multi-hop wireless networks. ExOR forwards each packet through a sequence of nodes, deferring the choice of each node in the sequence until after the previous node has transmitted the packet on its radio. ExOR then determines which node, of all the nodes that successfully received that transmission, is the node closest to the destination. That closest node transmits the packet. The result is that each hop moves the packet farther (on average) than the hops of the best possible predetermined route.

The ExOR design addresses the challenge of choosing a forwarding node after transmission using a distributed algorithm. First, when a node transmits a packet, it includes in the packet a simple schedule describing the priority order in which the potential receivers should forward the packet. The node computes the schedule based on shared measurements of inter-node delivery rates. ExOR then uses a distributed slotted MAC protocol for acknowledgments to ensure that the receivers agree who the highest priority receiver was.

The efficacy of ExOR depends mainly on the rate at which the reception probability falls off with distance. Simulations based on measured radio characteristics [6] suggest that ExOR reduces the total number of transmissions by nearly a factor of two over the best possible pre-determined route.

1 Introduction

Most unicast routing protocols choose a path of nodes through a network, and then send the data through that sequence of nodes. This approach makes the most sense when each pair of nodes is either linked by a wire, or else not linked, so that only linked nodes can communicate directly. If each link has a deterministic cost, then there will be one or more optimal routes between each pair. A routing protocol that first finds an optimal route, and then sends data along that specific route, is likely to perform well. Throughout this paper, we refer to the result of this strategy as the best possible predetermined route.

Multi-hop wireless networks deviate from the wired model in at least three ways. First, in principle all pairs of nodes can directly communicate over the radio, though perhaps with a high error rate. Second, nodes do not have to pick a particular target to send to (as if picking one of their links); at the radio level all packets are broadcast. Finally, the radio communication between a pair of nodes is not deterministic; packets arrive uncorrupted with some probability. That is, the properties that make pre-determined routes work well in wired networks do not hold in wireless networks.

This paper proposes a routing technique (ExOR) that takes advantage of the characteristics of wireless, rather than attempting to mask them. Instead of choosing a single route ahead of time, ExOR determines the path as the packet moves through the network, based on which nodes receive each transmission. This paper describes the details of ExOR, including a distributed MAC protocol that allows recipients to ensure that only one of them forwards the packet, and an algorithm that predicts which recipient is likely to be the most useful forwarder.

Simulations based on measured inter-node radio behavior suggest that ExOR will reduce the total number of transmissions required to forward a packet over a long multi-hop path by up to a factor of two. This improvement is somewhat sensitive to the node density (higher is better), to the length of the path (longer is better), and to the rate at which the packet reception probability falls off with distance (less steep is better). We expect that the technique could be made to work with 802.11b hardware with slight changes to the firmware which we describe in Section 3.5.

2 Intuition

In order to develop an intuition for why there might be room for improvement in multi-hop wireless routing, it is helpful to consider the simple network in Figure 1 in which there are a number of different possible routes from A to D. At one extreme, A could send directly to D in one hop, at the expense of sending each packet multiple times to recover from losses. At the other extreme, A could use the 3-hop route through B and C, at the expense of sending each packet multiple times since there are multiple hops.

Both of the above choices (and indeed any particular route)

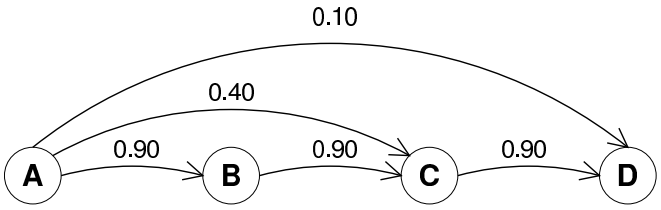

Figure 1: **Simple network example, with delivery ratios.**

leave some performance on the table. Node C and even D will hear many of the packets that A sends to B, and it is wasteful for B to forward such packets. If A tries to send directly to D, D may not receive the packet correctly but in many cases B or C will hear it, and it would be better for either of them to re-transmit the packet to D than for A to do so. The goal of ExOR is to take advantage of these opportunities to improve performance.

3 Details of the protocol

We begin with an overview of the protocol and supply details in the subsections below. The ExOR protocol consists of three stages: selecting the forwarding candidates, acknowledging transmissions, and deciding whether to forward a received packet. Throughout this description we assume each node in the network has a matrix containing an approximation of the loss rate for direct radio transmission between every pair of nodes. This matrix can be built using a link-state flooding scheme, in which nodes measure loss rates and periodically flood statistics updates.

The first node in an ExOR forwarding sequence chooses a candidate subset of all its neighboring nodes which could bring the packet closer to the destination. The sender lists this set in the packet header, prioritized by distance, as depicted in Figure 3. After transmission, each node that receives the packet looks for its address in the candidate list in the header. Each recipient then delays an amount of time determined by its position in the list before transmitting an acknowledgment. Each node looks at the set of acknowledgments it receives to decide whether it should forward the packet. The forwarding node rewrites the ExOR frame header with a new set of candidates and transmits the packet. This process is repeated until the ultimate destination receives the packet. The remainder of this section describes each phase in detail.

3.1 Selecting the candidate forwarder set

ExOR's performance is determined by its ability to choose a prioritized candidate set of nodes which bring a packet closest to its destination; for the networks described in this paper, simply choosing a candidate set based on the shortest number of hops (prioritized by delivery rates) results in good

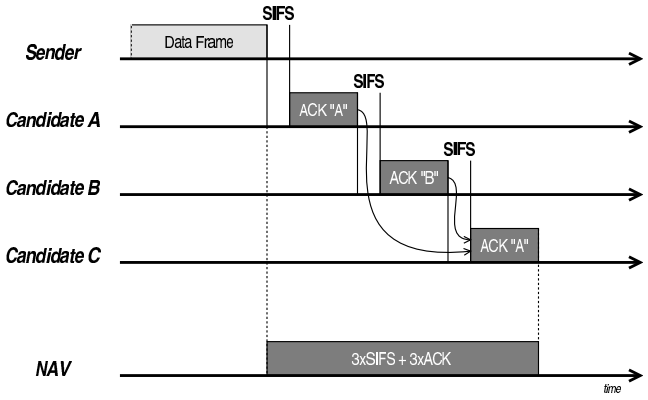

Figure 2: **Typical ExOR acknowledgment sequence.**

performance. ExOR chooses the prioritized candidate list as follows. It first identifies the shortest path to the destination, breaking ties between equally short paths using information from the delivery ratio matrix. The first node in this path is the highest priority candidate. Then ExOR deletes that node from the loss rate matrix, again finds the shortest route, and uses the first hop on that route as the candidate with second priority. It repeats this process to find the remaining candidates. The resulting candidate set for a given destination can be cached until the next update to the delivery ratio matrix. For example, a packet in the simple network from Figure 1 originating at A destined for C would have a candidate set of "D, C, B".

If this strategy were used in the cases of dense networks where ExOR only has a relatively small number of candidate selections, the candidate set would be filled with distant nodes with low reception probabilities. In such cases, a heuristic which chooses members from the complete range of next plausible hops should be used, ensuring the packet makes some progress on each transmission. Since the heuristic would likely depend on the characteristics of the network, we chose to study the simpler approach in this paper.

3.2 Acknowledgments

One of the major challenges of opportunistic routing is ensuring that the candidate nodes agree about which of them should forward the packet. We propose to use a modified version of the 802.11 MAC which reserves multiple slots of time for the receiving nodes to return acknowledgments. Instead of only indicating if the packet was successfully received, each acknowledgment contains the ID of the highest priority successful recipient known to the ACK's sender. All the candidates listen to all ACK slots before deciding whether to forward, in case a low-priority candidate's ACK reports a high-priority candidate's ID.

Including the ID of the sender of the highest-priority ACK

heard so far helps suppress duplicate forwarding. Suppose that node A hears a transmission, that A is the highest-priority candidate, and that A sends an ACK. Node B, the second-highest priority candidate, does not hear the ACK, but node C does hear the ACK. Suppose further that node B hears node C's ACK. If ACKs did not contain IDs, node B would forward the packet, since to its knowledge it is the highest-priority recipient. The fact that node C's ACK contains node A's ID indirectly notifies B that node A did receive the packet.

3.3 Deciding whether to forward a packet

Once the slotted acknowledgment window has passed, each candidate must make a local decision to forward or discard the packet. Only nodes which have not received acknowledgments containing the id of a higher priority candidate forward a packet. Occasionally multiple nodes will forward a packet due to acknowledgment reception failure. For this reason, each packet also contains a random nonce which forwarding nodes store in a cache to eliminate the possibility of forwarding the same packet multiple times. A packet is transmitted only if the nonce was not found in the cache.

3.4 Example Transmission

As an example of the ExOR routing protocol, consider the route taken by a packet in the simple network described in Figure 1 originating at node A and terminating at node D. In this case, the best candidate set for the packet is "D, C, B", as a successful transmission to D will deliver the packet to the destination, C will bring it very near and B will make some amount of progress in the correct direction and has a high probability of successful reception. Since there are a variety of reception possibilities, the particular case in which A transmits the packet and it is successfully received by nodes B and C, but not D serves as an example which exercises most of ExOR's mechanisms. After the initial transmission, the nodes now transmit acknowledgments in priority order, so the first acknowledgment slot belonging to node D is blank, the second slot from node C contains it's own node-id and the third slot corresponding with node B contains node C as well, since C has a higher priority within the candidate set. If all the nodes which received the packet successfully registered all the acknowledgments, we expect node C to become the new forwarding node. However, since there is a high probability that the acknowledgment from C to A was not received, node B's acknowledgment containing node C's node-id decreases the probability that A will retransmit the packet. Once node C has successfully determined itself as the responsible node, it forwards the packet.

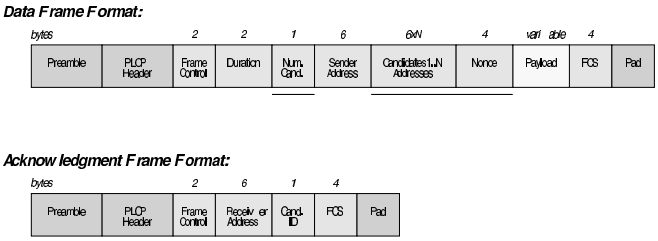

Figure 3: **ExOR Frame Formats. ExOR specific fields are underlined.**

3.5 802.11 Implementation Details

The data frame format and slotted acknowledgment mechanism required by ExOR can be implemented using 802.11 hardware [2] with minor MAC firmware modifications. Figure 3 shows the two modified MAC frame formats required by ExOR, in the context of the Physical Layer Convergence Protocol (PLCP) frame. The frame control field, which follows the 802.11 MAC specification, is set to either an otherwise reserved value, indicating the rest of the frame is either an ExOR data frame or an ExOR acknowledgment and not a standard data frame, control frame or management frame. The duration field of the data frames allocates time for acknowledgments from multiple candidates in the shared Network Allocation Vector (NAV). The 802.11 MAC uses the NAV for collision avoidance, so other nodes within radio range will not contend for the medium while it is reserved. The remainder of the data frame contains the candidate set in priority order, the packet's nonce, the payload, and the 802.11 Frame Check Sequence. The acknowledgment frames include the data frame sender's address and the position within the candidate list of the node with the highest priority acknowledgment. Since these acknowledgments must be transmitted in succession following a data frame (as shown in Figure 2), fine-grain timing similar to that used by standard acknowledgments is required.

ExOR adds a modest additional overhead to standard 802.11. Assuming a 802.11a physical layer operating at 6 Mbps with a SIFS time of 16 µs, preamble duration of 20 µs, PLCP header duration of 4 µs and 1500 byte frame payloads, the overhead of an ExOR data frame with four candidates and acknowledgments is 8.3% greater than unicast 802.11. The table below summarizes transmission times for each type of frame.

	Data Frame	ACK(s)	Total
802.11 unicast	2085 µs	58 µs	2143 µs
ExOR w/4 candidates	2092 µs	228 µs	2320 µs

4 Evaluation and Simulation Results

In order to determine the effectiveness of ExOR, we developed a simulation environment which models the operation of a multi-hop wireless network. The goal of the environment is to model a network with a large diameter in which we have perfect knowledge of delivery ratios in order to gain a better understanding of ExOR's potential when compared to the best possible pre-determined routes.

The simulated environment consists of nodes randomly placed in a plane, where the delivery ratio between two nodes is based on the distance-to-delivery relationship measured by Ganesan *et al.* [6] for Rene Motes using medium transmission power. The resulting distribution, displayed in Figure 4, is approximated by a linear function rather than an inverse-square or inverse-cube relation; understanding the underlying reason for this distribution is beyond the scope of this paper, but our experiments using an inverse-square curve indicate ExOR still performs well under such conditions. However, if the falloff is characterized by inverse-cube or greater, the benefit diminishes, as there are drastically fewer long-distance links available.

The simulator does not model medium contention or any radio propagation effects other than random packet loss based on a delivery ratio and assumes losses are not correlated with packet size. We also assume all the nodes are stationary throughout the simulation and thus delivery ratios do not change over time. A packet is transmitted by a single node a maximum of 8 times before being dropped.

As a baseline, we determined the best possible predetermined path by the perfect knowledge of link delivery ratios. In this case, we ran Dijkstra's shortest path algorithm with weights between the edges of $1/\text{fwdrate}\cdot\text{revrate}$, which determines the total number of transmissions required on average to send and acknowledge unicast packets along a given route. The ExOR protocol is implemented as described in the earlier section and is supplied with the same delivery rate matrix as the predetermined path algorithm.

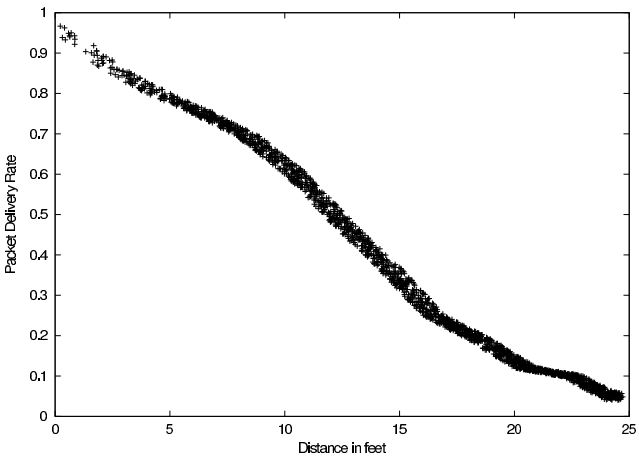

Figure 4: **Delivery Ratio versus Distance used in Simulation.**

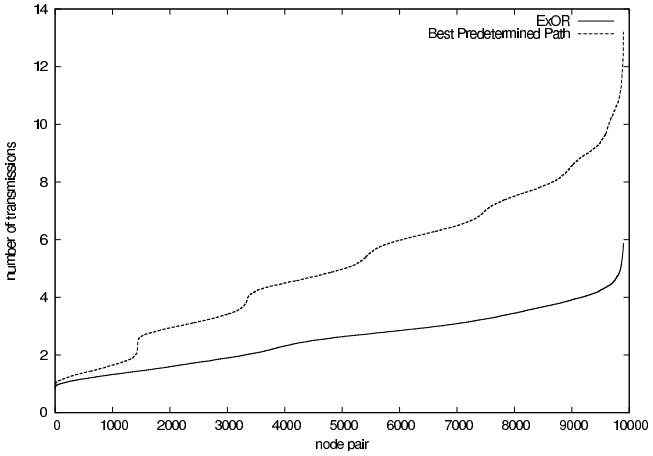

Figure 5: **Number of transmissions required to route a packet from source to destination, averaged over all pairs of nodes.**

4.1 Simulation Results

In this section we compare the performance of ExOR versus the best possible pre-determined route. We focus on two major aspects: total number of transmissions between all pairs and the distribution of links used by both approaches.

We simulate a network containing 100 nodes positioned randomly on a 50x50 sq ft plane and choose opportunistic routes using a candidate set size of 8. Figure 5 compares the average number of transmissions required to route 100 packets between every pair of nodes in the system. The results indicate opportunistic routing generally performs better than the best possible pre-determined route often by a margin of 55%. Simulations of longer networks in indicate improvements of up to 65% are feasible, which agrees with our intuition that the greatest benefit is then opportunistic routing has ample opportunity to skip intermediate hops. The sudden increases in number of transmissions by the best pre-determined route are an artifact of adding an extra hop into the route. The ExOR approach chooses routes opportunistically, so the slope remains smooth.

We next explore the distribution of successful transmission distances, as shown in Figure 6. In this experiment we consider 50 node pairs, choosing the five nodes closest to the edges of the network and plot the histogram of successful transmission distances when using the two approaches. We chose not to include measurements from all available node pairs, as a disproportionate numbers of pairs are within one or two transmissions of each other and therefore bias the his-

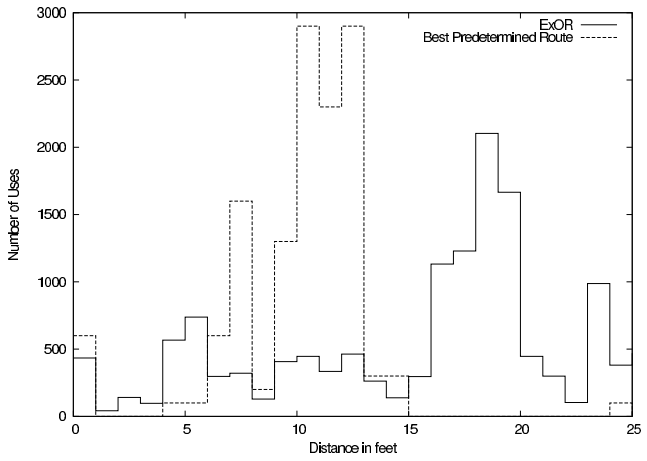

Figure 6: **Histogram of transmission distances when routing between 50 node pairs on either side of the network using ExOR and the best possible pre-determined routes.**

togram toward the smaller distances. The plot indicates opportunistic routing utilizes the longer transmission distances which are overlooked by the best predetermined routes and in general uses a wider variety of routes. Specifically, the bins to the right of those used by the predetermined routes indicate ExOR forwarded the packet farther that it would have with a statically determined route and those to the left indicate the acknowledgment came from a node a little farther along than the original sender, which indicates progress was made.

5 Related Work

ExOR is reminiscent of directed flooding protocols such as LAR [8]. ExOR combines much of the robustness of flooding with the efficiency of unicast paths, largely due to its ACK mechanism. Another difference is that while LAR uses geographic information to direct packets toward the destination, ExOR uses a link-state-style topology database.

Multi-path routing techniques, for example the braided multipath routing of Ganesan *et al.* [5], explore multiple routes in order to increase robustness or performance. Such techniques typically first choose multiple paths, and then use them either in parallel for performance or as primary/backup for reliability. ExOR uses paths that are similar to braided multi-paths, but uses all the nodes in parallel to increase each packet's robustness and delivery efficiency.

Geographic forwarding [3, 7] is opportunistic in the sense that a node decides the next hop when it forwards a packet, based on which neighbor is closest to the destination. The motivation for geographic forwarding is usually to reduce each node's routing state, rather than to increase the efficiency of the resulting paths; the lack of state makes it hard to choose globally high-quality routes. Blum *et al.* [1] improve local forwarding efficiency with a MAC technique similar to that in 3.2, though it depends on knowledge of geographic position and does not recover efficiently from loss of data or link-level ACK packets.

The slotted ACK mechanism was inspired by the randomized repair timers of the SRM multicast protocol [4].

ExOR can be viewed as a network-layer attempt to realize some of the benefits of radio-layer antenna and path diversity techniques [10, 9].

6 Conclusions and Future Work

In this paper, we presented ExOR, an opportunistic routing protocol which dynamically chooses paths on a per-transmission basis in a wireless network. In addition, we presented the algorithm and medium access layer changes necessary to efficiently implement the protocol. Results from simulations show opportunistic routing consistently outperforms the best predetermined routes in the number of transmissions with margins of up to 55%. We also found that ExOR takes advantage of longer transmission distances to forward the packet further on each hop than predetermined routes, which validates our intuition that transmissions are frequently received beyond intermediate nodes specified by even the best predetermined routes.

Future work on ExOR includes further analysis in simulation of node failure, candidate set sizes and heuristics, varied density, and alternative delivery ratio distributions. We also plan to implement ExOR on a hardware test-bed and are in the process of negotiating access to 802.11 MAC layer firmware which would allow us to implement the slotted acknowledgment window necessary for opportunistic routing.

References

[1] B. Blum, T. He, S. Son, and J. Stankovic. IGF: A state-free robust communication protocol for wireless sensor networks. Technical report CS-2003-11, University of Virginia CS Department, 2003.

[2] IEEE Computer Society LAN MAN Standards Committee. *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. New York, New York, 1997. IEEE Std. 802.11–1997.

[3] Gregory G. Finn. Routing and addressing problems in large metropolitan-scale internetworks. Technical report ISI/RR-87-180, USC ISI, March 1987.

[4] S. Floyd, V. Jacobson, C. Liu, and L. Zhang. A reliable multicast framework for light-weight sessions and application level framing. *IEEE/ACM Transactions on Networking*, 5(6):784–803, December 1997.

[5] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin. Highly-resilient, energy-efficient multipath routing in wireless sensor networks. *ACM Mobile Computing and Communications Review*, 5(4), October 2001.

[6] D. Ganesan, B. Krishnamachari, A. Woo, D. Culler, D. Estrin, and S. Wicker. Complex behavior at scale: An experimental study of low-power wireless sensor networks. Technical report UCLA/CSD-TR 02-0013, UCLA CS Department, 2002.

[7] Brad Karp and H. T. Kung. GPSR: Greedy perimeter stateless routing for wireless networks. In *Proc. ACM/IEEE MobiCom*, August 2000.

[8] Young-Bae Ko and Vaidya Nitin H. Location-Aided Routing (LAR) in mobile ad hoc networks. In *Proc. ACM/IEEE MobiCom*, pages 66–75, October 1998.

[9] J. N. Laneman and G. Wornell. Energy-efficient antenna sharing and relaying for wireless networks. In *IEEE Wireless Communications and Networking Conference*, September 2000.

[10] J. N. Laneman and G. Wornell. Exploiting distributed spatial diversity in wireless networks. In *Proc. Allerton Conference on Communications, Control, and Computing*, October 2000.