ROAR: A Multi-rate Opportunistic AODV Routing Protocol for Wireless Ad-Hoc Networks

Kwan-Wu Chin and Darryn Lowe

Wireless Technologies Laboratory
University of Wollongong
Northfields Avenue, NSW, Australia
{kwanwu, darrynl}@uow.edu.au

Abstract. In this paper, we outline a simple approach, called ROAR, that enables the Ad-Hoc On-Demand Distance Vector (AODV) routing protocol to strengthen its routes by recruiting neighbors of nodes on the least cost path as support nodes during the route construction process, and working closely with the medium access control (MAC) to employ an opportunistic forwarding scheme that takes advantage of the node diversity at each hop. We have implemented ROAR in the *ns-2* simulator over the IEEE 802.11a physical layer. From our simulation studies conducted using various network topologies and realistic radio propagation model, we find that ROAR increases AODV's packet delivery ratio and end-to-end throughput several orders of magnitude, in particular for hop count based routes. Therefore, ROAR provides a simple add-on that allows routing protocols to reap the benefits of diversity without relying on physical layer approaches.

1 Introduction

Wireless Ad-hoc network routing protocols must react quickly given challenges such as mobility and unreliable wireless links. In the recent past, test-bed experiences such as [2] and [3] have demonstrated the inadequacy of using hop count as a route metric. This is because routes are usually constructed using unreliable links that frequently experience deep fades since the hop count metric does not reflect link quality. As a result, routing protocols spend a significant amount of time reconstructing routes whenever one of the links on a route fades, or whenever a shorter but unreliable link becomes available.

A promising concept that can be used to combat the vagaries of the wireless channel is node or macro diversity. A common exploit in all diversity techniques is the uncorrelated or independent communication paths provided by nodes in the vicinity of the receiver. Therefore, when the receiver's link goes into a fade, it can rely on its neighbors to receive packets on its behalf. Thus far, the majority of the works that exploit macro diversity lie at the physical layer (see [10] for more details). This paper departs from these existing works in that we show how the network layer can take advantage of diversity at minimal cost using existing hardware. As we will see later, our solution is capable of achieving significant performance gains that otherwise would only be possible using multiple antennae or maximal ratio combining techniques at the physical layer.

T. Kunz and S.S. Ravi (Eds.): ADHOC-NOW 2006, LNCS 4104, pp. 87–100, 2006.

An important factor we need to take into consideration when forwarding packets is the underlying data rate adaptation algorithm. For example, the IEEE 802.11b [6] standard offers data rates ranging from 1 to 11 Mbps, and even higher for the IEEE 802.11a [7], which provides eight data rates to choose from, ranging from 6 to 54 Mbps. Therefore, a routing protocol must be cognizant of the MAC's multi-rate capability and also the behavior of the link adaptation algorithm being employed. For example, a receiver at d meters away may be able to support a packet error rate (PER) of 1% at 6 Mbps but quickly degrades to 10% PER if 54 Mbps is used. Thus, depending on the underlying link-adaptation algorithm's behavior, a routing protocol will observe different link quality. This is particularly important during the route construction process since routes are constructed using the base data rate whereas data are transmitted using much higher data rates.

Our approach called multi **R**ate **O**pportunistic **A**d-hoc on-demand distance vector (AODV) **R**outing protocol or ROAR provides a practical method to exploit node diversity and works closely with the underlying link-layer adaptation algorithm. ROAR extends the well-known AODV [13] routing protocol with the ability to set up support nodes passively along a given path. By exploiting the independent reception probability and link quality of each support node, ROAR is able to increase the packet delivery ratio and throughput of a given route. A key feature in ROAR is that we relax the constraint that packets must be forwarded to the next hop node identified in the routing table. Instead, any support node that belongs to the same relay zone can become the next-hop node. In other words, packets are forwarded between relay zones rather than hop-by-hop. Lastly, since support nodes at a given hop have the same routing information, packets are forwarded along the least cost path to the destination; a feature that avoids unnecessarily increasing the network's contention level.

We have implemented ROAR in the *ns-2* simulator and conducted extensive simulation studies over different network topologies. Our results show that ROAR offers a marked increase in packet delivery ratio and throughput. This is mainly due to the improved packet error rate (PER) at each hop, thus allowing higher data rates to be supported and less transient link failures that cause AODV to unnecessarily tear down routes. In our experiments, we observe several orders of magnitude increase in packet delivery ratio and throughput. Further, these performance gains are for routes constructed using the hop count metric. ROAR therefore provides a case to continue using hop count as a route metric in ad-hoc networks despite previous works documenting its negative effect on performance. Interested readers are referred to [4] for 18 reasons why ad-hoc networks should use long hop routes.

This paper has the following structure. We start by describing ROAR's motivation in Section 2. After that, in Section 3 we present ROAR and its integration with the underlying MAC and link-adaptation algorithm. We then present our simulation studies and results in Section 4. After that we outline relevant works in Section 5 before concluding in Section 6.

2 Motivation

Before presenting ROAR, we first quantify the benefits of node diversity; see Section 4 for simulation parameters. In the experiments to follow we use a simple point-to-point topology where initially there is only a sender and receiver nodes. The nodes then move apart by 10m after each simulation run. The traffic is constant bit rate (CBR) at 8 Mbps. After obtaining a baseline performance of a two nodes topology, we place additional support nodes in the vicinity of the receiver.

Figure 1 shows the benefits of having an additional node where the sender is able to sustain a given data rate for much longer distance. This is particularly significant for the base data rate, 6 Mbps, which see a 50 meter increase in transmission range before packet delivery ratio starts to decline.

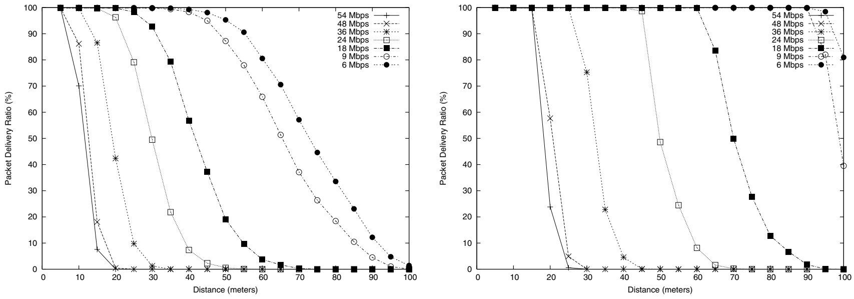

(a) No support node

(b) One support node

Fig. 1. The above figures show the benefits of having the receiver supported by an additional node

To determine the usefulness of diversity, an important question is the number of supports nodes required to yield a reasonable amount of performance gains. Figure 2 shows the benefit of adding up to six support nodes. It is clear that a limited number of nodes is sufficient to improve performance, however, as more nodes are added performance gains tend to become less.

Another investigation is to determine the benefits of node diversity on link-adaptation algorithms. Figure 3 shows the performance of two link-adaptation algorithms, AARF [9] and RBAR [5]. As expected, having more support nodes mean a given data rate is supported over a larger transmission range. Also, we see that AARF has a higher throughput than RBAR when the receiver is less than 50m since AARF does not incur the RTS/CTS exchange overhead. However, after 50m, RBAR becomes slightly superior due to data rate feedback from the receiver. Note, we restrict our comparison to these two algorithms since in [9] it has been shown that AARF has comparable performance to RBAR; minus the RTS/CTS overheads in addition to being deployable on current hardwares.

To summarize, we see that node diversity effectively strengthens a wireless link, as demonstrated by the increased packet delivery ratio and throughput for a given distance. We can therefore conclude that diversity will have benefits to any link adaptation algorithms. Apart from that, we see that the distance before a data rate's packet error rate worsen and the optimum data rate to use will depend on the diversity at the next hop. In the next section, we present a new protocol called ROAR that helps AODV reap the aforementioned performance gains.

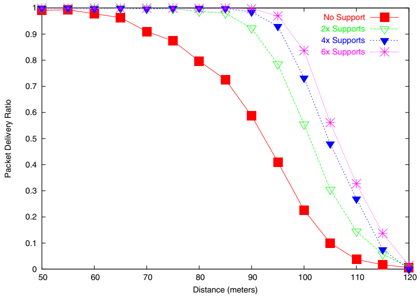

Fig. 2. The benefits of having more support nodes

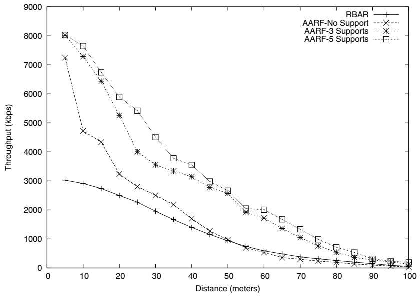

Fig. 3. AARF versus RBAR

3 ROAR: Protocol Description

3.1 AODV Background

AODV [13] is a reactive routing protocol. To construct a route, a route request (RREQ) is flooded across a MANET. Nodes with a route entry to the required destination send a route reply (RREP) message. To prevent unnecessary flooding of RREQ packets, nodes record the sequence number in each RREQ when they create a route entry to the RREQ originator. If a RREQ message with a lower or equal sequence number is received, it is discarded. Further, each node maintains a precursor list that, by recording the nodes using a given route, allows the determination of what neighboring nodes should receive a route error (RERR) message if the link to a given destination fails. In ROAR, we require each node record the number of duplicate RREQ messages transmitted by neighboring nodes, their IP addresses and link quality. We also augment both the RREQ and RREP messages to include the IP address of the previous and next hop node recorded in the routing table. For example in Figure 5, node-B would include node-C's IP address in the RREP before forwarding it to node-A.

We like to point out that ROAR can also be applied to other on-demand routing protocols, since it relies only on the route request/reply process; a function performed by all on-demand based routing protocols.

3.2 Relay Zone

A key concept in ROAR is *relay zones*. We like to point out that the term “zone” as applied here is different to conventional definition, which usually refers to a cluster of nodes managed by a leader and is used mainly to promote routing scalability. In ROAR, a relay “zone” is defined as follows.

Definition 1. *A zone is the set of nodes within the transmission range of one or more nodes that are on the least cost path to a given destination.*

That is to say, nodes within the transmission range of a next-hop node for a given destination are classified into the same zone. Figure 4 shows two zones where nodes in zone-B are within node-B’s reception area and similarly, nodes within node-A’s reception range belong to zone-A. Further, we see that node-3 belongs to two zones. It is worth pointing out that nodes A and B, referred to as *core* nodes, lie on the least cost path to a given destination and they are chosen during the route request and reply process. As will become clear later, no nodes play an active role in maintaining a zone. In fact, zone membership is defined as nodes that are listening to a multicast address constructed using the core node’s IP addresses and a prefix indicating whether its upstream or downstream, denoted as $MCAST_U$ and $MCAST_D$ respectively (see Section 3.2 for details).

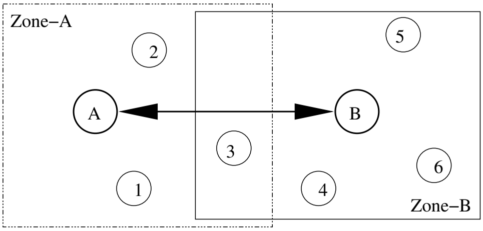

Fig. 4. Nodes within each zone play a support role to the node on the least cost path; nodes A and B

In the zone creation process we need to choose support nodes that will progress packets toward their intended destination quickly. This is because a support node may not have a good link to the next zone. In Figure 5, node-2 belongs to relay zones managed by node-B and node-A respectively. If node-A uses a high data rate when transmitting packets to node-B, it is likely that node-2 will receive a majority of the transmitted packets. However, notice that node-2 is much farther away to node-C (the next hop node) compared to nodes 3 to 6. Therefore, although node-2 can support a high data rate, it does not necessarily mean packets forwarded by it will make good progress toward their destination nor indicate it is in a good position to exploit node diversity in the next relay zone.

Relay Zone Construction. Relay zones are created once the node on the least cost path or core node is selected, which happens after receiving a RREP message. As the RREP

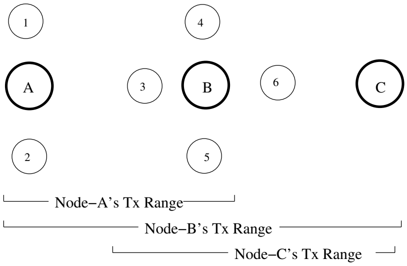

Fig. 5. Example used to construct relay zones

traverses from one node to the next, nodes overhearing the RREP message confirm their zone membership as follows.

Assume node-C has a route entry to the required target node and generates a RREP message to node-B. From Figure 5 we see nodes 3 to 6 are within node-C's transmission range, thus they will receive node-C's RREP message to node-B. Nodes 3 to 6 check to see whether they have created a route entry to the source due to an earlier RREQ message. If they have, the next step is determining whether they should join node-C or node-B's zone. The criteria for doing so must ensure node 3 or 6 has a good link to the next relay zone, thus ensuring packets are only forwarded on high-rate links.

In the discussion to follow we will define node i's next hop core node to a given destination or source as N_{i+1} or N_{i-1} respectively. We first consider setting up zones that ensure good progress toward the destination. To meet this requirement, support nodes must be located near or have good signal-to-noise ratio (SNR) to the core node; the node that sent the RREQ message previously. Once a neighboring node overhears the RREP message transmitted by the core node, it determines its SNR to the core node, and becomes a support node if its SNR is higher than a given threshold.

For packets going from the destination to the source, we used a slightly different procedure. Unlike before, we need to have support nodes close to node-B or N_{b-1} (node-A in this example). This is easily obtained since each RREQ contains the IP address of node N_{b-1}, therefore this information is already available when the RREQ came past. As stated above, nodes join zone-B if their SNR to node-B is higher than a predefined threshold or they can hear node N_{b-1}.

In the above cases, support nodes do not forward RREP messages. This ensures only core nodes forward the RREP message back to the source node. Each support node also updates their route entry to either node N_{i-1} or N_{i+1}'s IP address as the next hop to either the source or destination. This means support nodes have the same route entry for the given source and destination as the core node of the zone they belong to.

Route Maintenance. Route maintenance is trivial since we rely on AODV's RERR messages or route timeout. Alternatively, we could also rely on HELLO messages to indicate that a node is no longer a neighbor. In all cases, the routing layer removes the corresponding route information for a given neighbor and destinations using that neighbor as the next-hop node. Finally, for each destination, the routing layer informs the MAC to unsubscribe from the corresponding multicast address, thus removing itself from the relay zone. Note that, instead of unsubscribing from the relay zone, a support nodes within the relay may become the core node, hence enable quick local recovery from node or link failure. We leave this as future work.

Ethernet Multicast Address. As described earlier, each node in a given zone, including the core node, subscribes to an Ethernet multicast address. To construct a multicast address, we simply prepend the prefix $MCAST_U$ or $MCAST_D$, both 16-bits in length, to the hash of the core node and destination IP addresses. Notice that we have two sets of multicast address which we use to indicate support nodes that forward packets toward the source or destination node. This is done to ensure that only those nodes that contribute to the progress of packets join the zone. The prefix $MCAST_U$ and $MCAST_D$ denote Ethernet multicast address prefixes. The reason we need to hash the destination's IP address is to ensure that support nodes that are participating in other communication sessions do not intercept transmitted packets of other sessions since they may be not be within transmission range of a packet's next relay zone.

3.3 The MAC Layer

In the previous section we have shown how relay zones are created and also defined what it means to belong to a given zone. In this section, we show how the IEEE 802.11 MAC makes use of relay zones.

Before describing our modifications, we first highlight some key design criteria.

1. *Minimal modifications to the IEEE 802.11b standard.* This criterion ensures that ROAR is practical and is deployable using current hardware. As we will see later, ROAR's only modification consists of introducing a slight random delay in the clear channel assessment process to avoid acknowledgment message collision.
2. *Minimal probing.* We want to avoid having to obtain data rate estimates from nodes in the next relay zone, since this incurs non-negligible delay and signaling overheads, in addition to requiring changes to IEEE 802.11's RTS and CTS exchange.

Data Packet Transmission. The forwarding process at the network layer proceeds as normal where the routing daemon obtains the next-hop or core node's IP address from the routing table which is then passed to the MAC to construct an Ethernet multicast address. In the descriptions to follow we assume all control messages are transmitted at the base rate. Further, we will use Figure 6 to illustrate the steps involved.

The packet transmission proceeds as usual. A node senses the channel, and if the channel is free, it transmits the packet using the next relay zone's multicast address. The data rate used will be dependent on the packet transmission history to the given neighbor. In this example, since there are three nodes listening to the multicast, we need to resolve which one of them will send an acknowledgment message to the sender. To resolve this issue, each node performs the clear channel assessment operation for a small random period. At the end of the period, if the channel is free, it transmits the acknowledgment message back to the sender. The packet is then forwarded to the routing layer. Otherwise, if another transmission is heard, it discards the data packet.

Link Adaptation Algorithm. The multi-rate capability of current wireless systems has spurred researchers to develop link adaptation algorithms. For example, OAR [15]

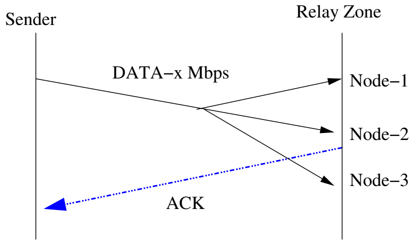

Fig. 6. Data transmission

determines the best data rate to use by probing the channel using a request-to-send (RTS) message and having the receiver return its estimated data rate in the returning clear-to-send (CTS) message. Interestingly, in [1] and also in [9] the authors commented that, based on their test-bed experience, estimating data rate using instantaneous SNR is unreliable, thus giving more credence to opportunistic forwarding schemes.

Due our aforementioned criteria and the above reason, we have chosen to use AARF [9], thus avoid the use of the RTS/CTS exchange. AARF's key advantages is its simplicity and zero signaling overheads, which is particularly important when dealing with a large number of support nodes. Further, it does not require any modifications to the RTS/CTS handshake nor changes to packet formats. It has also been shown to have comparable performance to RBAR [5]; a scheme that relies on the exchange of RTS and CTS messages in order to obtain a suitable data rate estimate before transmission. Lastly, it can be deployed immediately without any modifications to the MAC.

AARF works as follows [9]:

- After M successful transmissions, the sender increases its data rate to the next level. Initially, M is set to 10.
- Define the first packet transmitted using the new data rate as the probe packet. If the probe packet fails at the new data rate, reduce data rate immediately and double the value of M. Note, M has a maximum value of 50.

AARF has an important characteristic. That is, the ability to settle on an M value suited for a given path loss. This optimization is particular important in the context of routing in wireless ad-hoc networks because increasing the link rate aimlessly may result in a higher packet loss, thus causing the routing protocol to incorrectly conclude that the link to a neighbor is unstable.

4 Experiments and Results

This section presents our experiments and results. We equip *ns-2* (v2.28)'s [12] with IEEE 802.11a [7], thus giving us eight data rates; 6 to 54 Mbps. To ensure a realistic channel model, we used the Ricean propagation model modifications of [15] and fixed the velocity value to 2.5. Unless indicated otherwise, all our experiments are performed with $K = 0$. Further, a new seed is used to initialize the random number generator for each simulation run. Lastly, we extended AODV with ETX [3]. Note, for simulation involving ETX, we allow each node to average up to 600 ETX values before starting the route construction process.

4.1 Linear Topology

We start our experimentation with a linear topology, see Figure 7, to first measure the performance gains in a simple setting before moving onto more general topologies in Section 4.2. In this experiment, node-A sends a CBR stream at 10 KB/s to node-E. We increase the distance between nodes by five meters after 50 runs and record the average packet delivery ratio (PDR) and throughput. In addition, depending on the experiment, we add one to three support nodes in the vicinity of nodes B, C and D. In our simulation, AODV's local repair feature is disabled.

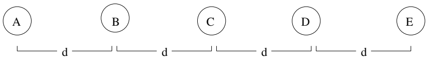

Fig. 7. Linear topology. Additional support nodes are added near the vicinity of nodes B, C and D.

Figure 8(a) shows the PDR when AODV uses hop count as a route metric with three support nodes. In this experiment we have omitted ETX since the clustering of nodes meant that there are no alternative reliable paths that can be constructed using ETX. An important aspect of ROAR is improving the link when it is unreliable. For example, see Figure 8(b), as nodes become more distant from each other, i.e., PER gradually becomes higher, ROAR is able to increase a route's PDR significantly. This is important for hop based routes since they tend to select unreliable next hop node, therefore, by using ROAR the benefits hop count are maintained; with the negative effects due to wireless fading minimized.

Figure 9 shows the average throughput for the linear topology. Due to the lower PER, AARF has a better chance of using a higher data rate. Ensuring a link's PER is low is important. This is because AARF reduces the current data rate whenever the MAC has given up retransmitting a data packet (7-retries) or when the probe packet is lost. Another benefit is that AARF is able to reach steady state sooner, recalling that the IEEE 802.11a physical layer supports eight data rates and AARF only increases the data rate after 10 consecutive successful transmissions. In other words, since there is less packet loss, AARF is able to quickly settle on a data rate commensurate with link condition.

4.2 Node Density

In this experiment we investigate the effect of node density. We use a grid size of 200x50 meters. We fix the position of the sender and receiver coordinate $< 10, 10 >$ and $< 190, 90 >$ respectively. We then increase the number of nodes from 10 to 60. In this experiment, we set the Ricean propagation model's K value to three. Also, in each simulation, the sender generates a CBR traffic with a maximum of 2000 packets. We record the PDR of 50 simulation runs, and plot the median PDR for each network density. The reason we plot the median instead of the mean is due to the variability caused by wireless errors and varying path loss from random node positions.

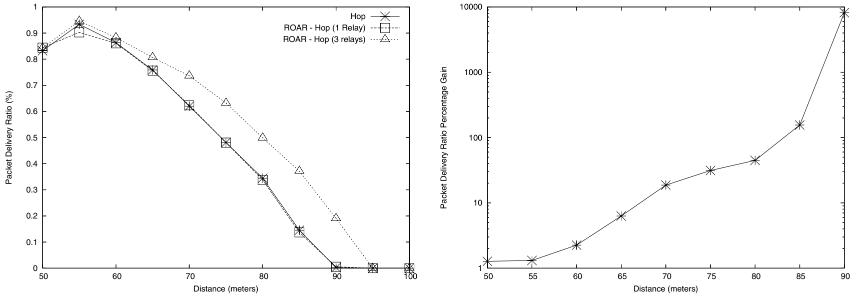

(a) Average PDR (b) PDR gain.

Fig. 8. Figure (a) shows the PDR for the linear topology, and (b) is the corresponding PDR gain for ROAR with three support nodes over a scenario with no support nodes

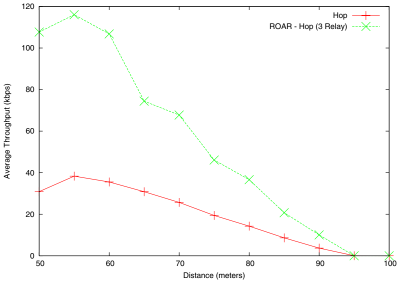

Fig. 9. Average throughput for linear topology

Figure 10 shows the performance of AODV with and without ROAR for two route metrics. For both metrics, their PDR increased up to 25% compared to AODV without ROAR. As network density rise, most of the links on a route will be within transmission range and in close proximity. Further, they have neighbors close by which ROAR can exploit. All these mean that wireless losses become less due to more reliable links on a given route, thus leading to a gradual performance improvement as node density rises. Similar to PDR, ROAR has a significant impact on throughput, as illustrated in Figure 11. Observe that as node density rises, the performance of AODV using ETX becomes lower due to increasing path length. However, in both ETX and hop count cases, ROAR sees up to 50% increase in throughput over basic AODV.

4.3 Recruitment Threshold

In this experiment we adjust the data rate thresholds used to recruit support nodes. Threshold refers to the minimum data rate a neighbor must support to qualify as a support node. For example, a threshold of 50 means only those neighbors that can support a data rate of 54 Mbps, whereas a threshold of 10 means those neighbors that can support a data rate of at least 10 Mbps, i.e., 12 to 54 Mbps. In this experiment we place

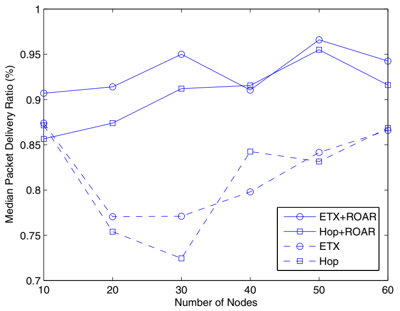

Fig. 10. Effect of node density on packet delivery ratio

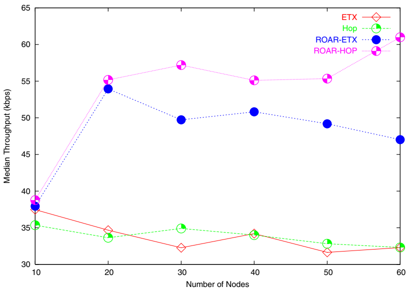

Fig. 11. Effect of node density on throughput

50 nodes randomly in a grid size of 200x50 meters, and calculate the average packet delivery ratio for 50 simulation runs.

Figure 12 shows the effect of using different thresholds in two different network densities; 20 and 40 nodes. We see that for both network densities, having a threshold at 40, i.e., those neighbors capable of supporting 48 and 54 Mbps, is sufficient to obtain a high PDR. At the highest threshold of 10 we see that PDR drops dramatically. This is due to increased contention resulting from duplicate packets since support nodes are out of each other's transmission range, thereby unable to sense that the other has transmitted an acknowledgment for a given data packet. As a result, packets are dropped, followed by the MAC informing AODV that the link to a neighbor is broken. This then causes the removal of packets buffered on the said link and re-initiation of the route construction process.

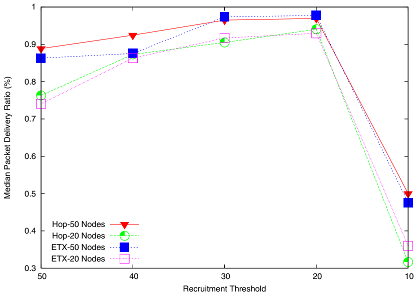

Fig. 12. The effect of support node recruitment threshold on PDR

5 Related Work

Researchers are beginning to investigate new packet forwarding paradigms that exploit node diversity to combat the vagaries of the wireless channel. Recently Biswas et al. [1] presented a protocol where the sender transmits a batch of packets without first having to probe the channel condition nor the need to specifically target a set of relays that will forward the packets onward. Packets that failed to reach the receiver are retransmitted by relays in priority order set by the sender. Further, these relays are required to track which packets have been received by the destination to minimize redundant transmissions. In a different work, Yu et al. [18] proposed that a neighboring node of the receiver send an acknowledgment message to the sender if the receiver failed to receive the packet. ROAR has similar objectives to both of the aforementioned works. However, ROAR requires less state exchanges and coordination between forwarding nodes. Moreover, we consider the dynamics of the underlying link adaptation algorithm. Finally, packets are forwarded using the shortest hop path, thus it has the advantage of not adding to the contention level.

ROAR is also closely related to two other works. In [17], the authors use two hops neighborhood information to select relay nodes with good link quality. On the other hand, [8] relies on a multicast routing protocol to obtain candidate relays' addresses. In both cases, the relay addresses are included in RTS messages to determine the relay with the best channel condition. The authors of [8] also suggested using other information such as queue length and geographical distance of neighboring nodes. The list of addresses embedded in the RTS message are then used by the relays to determine which node has the highest priority to transmit a CTS message. We improve upon these works as follow. ROAR eavesdrop on AODV's RREQ and RREP messages, thus relays are recruited passively and they are identified using a Ethernet multicast address. Secondly, a sender does not need to know relays' IP addresses since ROAR uses an opportunistic forwarding paradigm that does not require it to determine the best relay or support node to use. Finally, ROAR considers the effects of link adaptation.

The zone creation aspect of ROAR resembles AODV-BR [11]. AODV-BR establishes backup nodes at each hop where upon a route failure, they are called upon to provide an alternative path to the next node on the shortest hop path. AODV-BR also relies on RREP messages to set up backup nodes. ROAR on the other hand establishes so called backup nodes for a different reason, that is to increase node diversity which the MAC can then exploit. Another difference is that ROAR allows support nodes in different relay zones to communicate directly. However, in AODV-BR, a backup node's next hop node is always a node on the shortest hop path.

There have been quite a number of works related to constructing multiple paths in MANETs to increase reliability [16], limit packet loss and load balancing [14]. Although the end goals are similar, our works defer in that we are not concern with end-to-end disjoint paths nor require a multicast routing protocol.

6 Conclusion

ROAR incorporates AODV with two interesting concepts: node diversity and opportunistic forwarding. These concepts enable ROAR to exploit the vagaries of the wireless channel and to do so with minimal signaling overheads. Furthermore, these concepts are realized with minimal modifications to AODV and the IEEE 802.11 MAC.

Through extensive simulation studies, we evaluated ROAR in various network topologies and found that ROAR provides orders of magnitude performance gain at minimal cost. Apart from that, we show that ROAR strengthens routes constructed using hop count. Finally, ROAR provides a simple network layer solution that offers some of the performance gains that have been touted by cooperative diversity techniques at the physical layer.

Acknowledgment

This project is partially funded by ARC Discovery Project DP0559769.

References

1. S. Biswas and R. Morris. ExOR: Opportunistic multi-hop routing for wireless networks. In *ACM SIGCOMM'2005*, Pennsylvania, USA, Sept. 2005.
2. K.-W. Chin, J. Judge, A. Williams, and R. Kermode. Implementation experience with MANET routing protocols. *ACM/SIGCOMM Computer Communications Review*, 32(5), Nov. 2002.
3. D. S. J. DeCouto, D. Aguayo, J. Bicket, and R. Morris. A high throughput path metric for multi-hop wireless routing. In *MOBICOM'03*, San Diego, USA, Sept. 2003.
4. M. Haenggi and D. Puccinelli. Routing in ad-hoc networks: A case for long hops. *IEEE Communicaitons Magazine*, 44(10):93–101, Oct. 2005.
5. G. Holland, N. Vaidya, and P. Bahl. A rate-adaptive MAC protocol for multi-hop wireless networks. In *ACM MOBICOM'2001*, Rome, Italy, Oct. 2001.
6. IEEE. Wireless LAN medium access control (MAC) and physical layer (PHY) specifications. IEEE Standard, Draft 802.11, 1997.
7. IEEE. Part 11a: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications: High speed physical layer in the 5 GHz band, standard specification. IEEE Std 802.11a-1999, 1999.
8. S. Jain and S. R. Das. Exploiting path diversity in the link-layer in wireless ad-hoc networks. In *Proceedings of IEEE WoWMoM*, Taormina, Italy, June 2005.
9. M. Lacage, M. Manshaei, and T. Turletti. IEEE 802.11 rate adaptation: A practical approach. In *ACM MSWiM'2004*, Venice, Italy, Oct. 2004.
10. J. Laneman, D. Tse, and G. Wornell. Cooperative diversity in wireless networks: Efficient protocols and outage behaviour. *IEEE Transactions on Information Theory*, 50(12):3062–3080, Dec. 2004.
11. S.-J. Lee and M. Gerla. AODV-BR: backup routing in ad-hoc networks. In *IEEE Wireless Communications and Networking Conference (WCNC'00)*, Chicago, IL, Sept. 2000.
12. S. McCanne and S. Floyd. ns network simulator-2. http://www.isi.edu/nsname/ns/.
13. C. E. Perkins, E. M. Royer, and S. R. Das. Ad hoc on-demand distance vector (AODV) routing. draft-ietf-manet-aodv-06.txt, July 2000.
14. M. Perlman, Z. J. Haas, P. Sholander, and S. Tabrizi. On the impact of alternate path routing for load balancing in mobile ad-hoc networks. In *IEEE/ACM MobiHOC'2000*, Chicago,IL, Aug. 2000.
15. B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly. Opportunistic media access for multirate ad-hoc networks. In *ACM MOBICOM*, Atlanta, Georgia, USA, Sept. 2002.

16. A. Tsirigos and Z. Haas. Analysis of multipath routing - part ii: Mitigation of the effects of frequently changing network topologies. *IEEE Transactions on Wireless Communicaitons*, 3(2):138–146, Mar. 2004.
17. J. Wang, H. Zhai, W. Liu, and Y. Fang. Reliable and efficient packet forwarding by utilizing path diversity in wireless ad-hoc networks. In *Proceedings IEEE Milcom'04*, 2004.
18. C. Yu, K. G. Shin, and L. Song. Link-layer salvaging for making routing progress in mobile ad-hoc networks. In *ACM MobiHOC*, Urbana-Champaign, Illinois, USA, May 2005.