

MAC-Layer Anycasting in Ad Hoc Networks

Romit Roy Choudhury⁺ and Nitin H. Vaidya*

⁺Dept. of Computer Science,

*Dept. of Electrical and Computer Engineering, and
Coordinated Science Laboratory,
University of Illinois at Urbana-Champaign

Abstract—A wireless ad hoc network is formed by a group of wireless hosts, without the use of any infrastructure. To enable communication, hosts cooperate among themselves to forward packets on behalf of each other. A key challenge in ad hoc networks lies in designing efficient routing strategies. While several routing protocols have been proposed, most of them aim to select one optimal route between the source and destination. The MAC layer at each intermediate node is then required to forward packets to the next downstream node on that route. We argue that choosing a single optimal route at the network layer may not be sufficient. Knowledge of short-term channel conditions at the MAC layer can play an important role in improving end-to-end performance. Instantaneous interference, channel contention, power constraints and other considerations may be taken into account along with the network layer’s long-term view. This paper proposes *MAC-layer anycasting* – a forwarding strategy that combines the guidelines from the network layer, with MAC layer knowledge of the local channel. We describe some applications of *MAC-layer anycasting*, and discuss the performance related tradeoffs.

I. INTRODUCTION

Wireless Ad hoc networks are infrastructureless multi-hop networks in which nodes behave as mobile routers. Routing protocols attempt to choose “optimal” routes based on some optimality criteria (e.g., number of hops). However, in the process of selecting an optimal route, the routing protocol is often faced with the decision to choose between two equally good routes. Ties are often broken randomly. MAC-layer anycasting is a proposal that aims to utilize the knowledge of instantaneous channel condition in selecting the suitable downstream neighbor on shorter time scales. The observation that routes chosen by the network layer are “optimal” on a longer time scale, and ignores the possibility of transient variations in link conditions, motivates our work on MAC-layer anycasting.

The key idea behind MAC-layer anycasting is to achieve the goals of the network layer, while invoking short-term optimizations at the MAC layer, based on local channel conditions. With the proposed approach, the network layer is given the option of specifying multiple downstream destinations to the MAC protocol. The MAC protocol assumes that forwarding the packet to any one of these destinations is acceptable to the routing layer. Depending on the current channel state, the MAC layer then forwards the packet to one of the specified neighbors. Out-of-order packet delivery is a potential problem with proposed anycasting. We discuss this, and other tradeoffs associated with anycasting, later in the paper.

This work is supported in part by National Science Foundation (NSF) under grants 01-96410 and 01-25859.

II. PRELIMINARIES

Routing protocols can be broadly classified into “source-routed” or “table-driven” protocols [1],[2]. In source routing [3], the sender of a packet completely specifies the route that the packet must traverse to reach its final destination. Johnson *et al.* proposed dynamic source routing (DSR) [3] in which the sender node floods a route request (RREQ) probe in search of a route to the destination. Intermediate nodes that forward this request probe, append their identifiers to the probe. The probe that arrives first at the destination is assumed to have arrived on the optimal path. DSR uses this path for subsequent communication.

Table-driven routing protocols store routing information locally [4],[5],[6][7]. Nodes exchange routing messages, either reactively or periodically, to update each other about the status of links in the network. When a node intends to send data packets to another node, it consults its routing tables for a route to the destination. It forwards the data packet to the appropriate neighbor in the route, who in turn consults its own tables to forward the packet further. An intermediate node is often faced with the decision to choose between two of its neighbors, both of which may be equally good for forwarding the packet to the final destination. Ties are broken randomly, without respecting the possibility that one of the nodes may not be suitable for immediate transmission. We believe that anycasting can be useful here – the MAC layer can make educated decisions in such scenarios, leading to potential benefits in performance. In this paper, we would refer to table driven protocols while discussing the details of MAC-layer anycasting. Issues arising from the use of source routing will be discussed separately in Section V.

Roy *et al.* propose the notion of *maximally zone disjoint* routes [6]. Based on previous traffic conditions, a sender selects routes that can maximally bypass congested regions. Our idea of anycasting differs from [6] in the sense that we base our forwarding decisions on factors that vary on a shorter time scale. The routing layer only provides a set of acceptable options (not all of which may be optimal). The MAC layer then chooses the next hop depending on the instantaneous network condition. Pursley *et al.* [8] proposed the idea of using “decoder side information” to aid forwarding decisions. By observing the number of correct symbols received (from a sequence of known transmitted symbols), the receiver may be able to estimate, statistically, the reliability of the link. The authors propose a metric,

resistance, which is indicative of link quality. Using this metric, a node examines two outgoing links, and transmits the packet over the one with lower *resistance*. While this scheme handles variation in channel fluctuations, it does not consider issues related to the MAC layer. MAC-layer anycasting adapts to several MAC protocol constraints, as detailed in the rest of the paper.

Larsson [9] presents the idea of “selection diversity forwarding”, in which a transmitter includes a multicast address (or a list of addresses) in the data packet. Neighbors of the node that are included in the multicast group (or the address list), reply to the packet serially with an ACK packet. The transmitter chooses one of its neighbors, based on the guidelines of the routing layer and the current link conditions learned from the data-ACK exchange. A “forwarding order” is now transmitted to the chosen neighbor, requiring it to forward the packet further. The chosen neighbor replies to the “forwarding order” with a “forwarding order ACK”. Clearly, waiting for all the replies before initiating the “forwarding order” may be wasteful. Jain *et al.* [10] propose an improvement on the protocol in [9]. The authors propose to specify the list of addresses (similar to [9]) in order of priority. The protocol requires all nodes, included in the address list, to reply in sequence of priority, with the highest priority first. Upon receiving the first reply (not always from the highest priority node), the transmitter immediately begins data packet transmission to that node. This reduces the overhead associated with waiting for multiple replies before transmitting a packet. Unlike [9], the order of priority must be specified *a priori* without knowledge of the instantaneous link conditions. In addition, specifying preferences and multiple addresses increases packet-size, leading to higher control overhead.

Although similar in spirit, MAC-layer anycasting can be distinguished from the body of existing work. The key distinction lies in the basis of decision-making. Observe that most of the previous schemes rely on probing the channel in some form, and choose the suitable neighbor based on explicit or implicit feedbacks. We argue that in several cases, waiting for feedbacks may not be necessary – the MAC layer may already possess necessary information. For example (more examples elaborated later), we observe that the MAC layer may be aware of permissible transmit power-levels at a given point of time. Previous schemes may probe the channel with an impermissible power level, obtain a negative feedback, and converge to the permissible power level. Clearly, using the knowledge available at the MAC layer can be useful in such scenarios. We discuss some wireless medium access control (MAC) protocols next.

We assume that the reader is familiar with the IEEE 802.11 protocol [11]. Briefly, when using 802.11, an exchange of request to send(RTS)/clear to send(CTS) precedes DATA communication. Nodes that overhear the RTS/CTS defer their own transmissions, for the proposed duration of the DATA communication. Once the DATA packet has been transmitted, the receiver replies with an ACK to acknowledge successful reception.

Several proposals in the recent past have tuned 802.11. However, the key idea of the protocol remains unchanged. Re-

cently, with advances in antenna technology, several protocols have been proposed that use directional antennas at the MAC layer¹ [12][13],[14],[15],[16]. The key ideas when using directional antennas may be summarized as follows. Due to the ability to transmit signals in a desired direction, most of the protocols propose to use a combination of directional and omnidirectional RTS/CTS/DATA and/or ACK. Spatial reuse of the channel increases due to reduced interference. The notion of directional NAV [15],[16] enables a node to initiate transmissions that will not interfere with ongoing communication. Range extension, possible due to the higher gain of antenna beams, is an additional benefit – fewer-hop routes can be formed between the source and the destination [14],[16]. Although promising, directional antennas also pose some difficulties. Neighbor discovery [13],[14], new types of hidden terminals [16], deafness [16],[17] are some of the problems that arise from directional communication. We believe that anycasting can help, when using directional antennas.

Research on multi-user diversity in medium access control protocols has also been a topic of interest. Qin *et al.* [18] proposes a *channel-aware ALOHA* protocol, that schedules transmissions based on instantaneous channel conditions. Using a distributed approach, the protocol requires a node to transmit when its local channel conditions are favorable. Tsatsanis *et al.* [19] proposed “network assisted diversity protocols”, where the possibility of exploiting corrupted packets has been explored. Put differently, the authors propose the idea of allowing multiple transmitters to collide multiple times (synchronously). From the vector of corrupted packets, the receiver then separates the individual packets, using known signal processing algorithms. DeCouto *et al.* [20] have recently proposed an ETX metric to favor paths that are characterized by fewer losses and retransmissions. Put differently, while making the routing decisions, the network layer considers the information available at the MAC layer. However, once a route has been chosen, it is used irrespective of the possible changes in instantaneous channel conditions. Yarvis *et al.* [21] have also proposed similar ideas in the context of sensor networks. The key idea in this paper is somewhat opposite to that of [20], [21]. The MAC layer requires the network layer to supply a set of routes, that it deems suitable. Unlike the above approaches, the MAC layer performs the final decision of choosing the neighbor that appears to be most appropriate at that instant of time.

III. MAC-LAYER ANYCASTING

MAC-layer anycasting can be envisioned as an enhancement to existing MAC and routing protocols. In the rest of this paper, we would call a routing protocol “basic” if it has not been “enhanced” with the anycasting features. One possible architecture to implement MAC-layer anycasting is shown in Figure 1. This section discusses the framework of MAC layer anycasting, in the context of a generic MAC and routing protocol. We also propose a simple variation, named *Ordered Anycasting*². Later,

¹Although the 802.11 protocol can be used over directional antennas, performance improvements may not be achieved.

²We first present this concept in a proposal submitted to NSF in December 2001

we visit the applications of anycasting and discuss the tradeoffs in the context of wireless ad hoc networks.




Fig. 1. An Anycasting Framework.

The anycast framework requires the “basic” routing protocol to discover/maintain multiple routes for each flow, whenever possible. Clearly, all the discovered routes may not be equally good. When a packet arrives at the network layer, the routing protocol consults the routing state to determine the routes that may be available for the packet’s final destination. From these available routes, the routing protocol selects a subset containing K routes that may be deemed as the best. The network layer now forms what we call the *anycast group*. The anycast group contains the set of distinct next-hop neighbors, on the selected K routes. As an example, for a packet destined to D, the anycast group specified by the network layer at node S in Figure 2 could be the set (A,X). The packet and the anycast group are then handed down to the MAC layer. Upon receiving the packet, and the anycast group, the MAC layer must select any one suitable neighbor and attempt transmission to it. Instantaneous network conditions may play an important role in determining the selection. The next section presents some of the potential applications of anycasting, and illustrates how the neighbor selection policies may be designed. However, first we propose a simple variation to anycasting, named ordered anycasting.

Ordered anycasting

The routing layer at a node may discover multiple routes to a particular destination. All the routes may not be optimal. For example, if routes R1 and R2 are equally good (e.g., in terms of hop-count), and if both are better than route R3, then the network layer may desire to use R3, only if communication over routes R1 or R2 is currently not possible. Ordered anycasting is a simple variation to anycasting that aims to achieve exactly this. The routing layer ranks the members of the anycast group in order of its preference. The MAC layer attempts communication to a node, only if all other nodes higher in the preference order, have proved to be “unavailable”.

IV. APPLICATIONS OF ANYCASTING

This section discusses the shortcomings of existing MAC and routing protocols. We investigate example scenarios to understand consequences of these shortcomings, and show how anycasting can be useful.

A. MAC constraints

Consider the scenario in Figure 2. Assume that the network layer at node S selects a route to node D through intermedi-

ate nodes A and B (i.e., S,A,B,D). Also, assume that over the lifetime of the route, other flows in the network forward packets through node E. Clearly, using a MAC protocol like IEEE 802.11, node E would require its neighbors to remain silent while it is receiving packets from node F. Node A, in the neighborhood of E, must therefore refrain from communication, while node E has reserved the channel. If node S transmits an RTS to A during this interval, A would be unable to reply with a CTS. Node S will interpret the absence of a CTS as a sign of collision at A, and backoff exponentially before reattempting transmission³. In the meantime, E might reserve the channel for yet another communication. S would continue to retransmit, until A is available for communication. Clearly, forwarding data packets on this route gets delayed due to repeated failures.




Fig. 2. An example scenario illustrating the possibility of anycasting

Observe from Figure 2 that alternate routes exist between nodes S and D – for example $\{S,X,B,D\}$. Node X need not refrain from communication when E has reserved the channel, and can therefore be a potential candidate for forwarding downstream packets. Hop-count remains the same if either of the routes, $\{S,A,B,D\}$ or $\{S,X,B,D\}$ is used. By forwarding the packet to X, node S can avoid the possibility of multiple retransmissions on link S-A. End-to-end delay reduces if X is instantaneously available for communication. Similar optimizations can be invoked at all the intermediate nodes on the route. Transmitting each packet to *any one* from an acceptable set of next-hop neighbors can be achieved using MAC-layer anycasting.

The above problem might be more pronounced when using directional antennas in ad hoc networks. Consider Figure 3, where node A is engaged in communication with node B. Since node A is beamformed in the direction of A, it would be unable to receive signals from S. If MAC-layer anycasting is not used, the network layer would have specified node A as the next-hop for the packet, and hence, node S would continue to attempt retransmission to A without success. If A has multiple packets to send to B, link S-A can be unavailable for a long duration. Performance can degrade significantly, as evaluated in detail in [23]. With MAC-layer anycasting, node S can exploit the option of forwarding the packet to node X. X may be able to respond immediately to S, even if A is busy. Unnecessary retransmissions can be avoided, leading to lower delay and fewer packet

³This problem arises in several scenarios in wireless medium access control. The impact of this problem has been evaluated in [16],[22].

drops.

Link unavailability may also arise if nodes are programmed to periodically turn off their transceivers, to conserve energy [24],[25]. In such scenarios, a link between nodes S and A would be unavailable if A has powered itself off (i.e., sleeping). While some proposals have addressed the problem of sleeping, anycasting could also be a useful solution, especially when the network is dense.

If link unavailability is the dominating motivation to implement anycasting, the neighbor selection policy must be designed accordingly. We propose one possible design, named *instantaneous link probing*. *Instantaneous link probing* aims to reduce the impact of link unavailability, by trying to communicate to each of the members in the anycast group. Using this mechanism, the MAC protocol selects next-hop neighbors in a round robin manner, and attempts transmission to each of them. The order of choosing the next-hop neighbors, and the number of retransmissions to each of them, can be a function of the neighbor's recent behavior. Next-hop neighbors that have responded quickly in the past can be attempted earlier than those that have recently been unavailable. In addition, more retransmissions can be allocated to the former. For an example, if the anycast group is specified as (X,A), then one possibility for node S could be to attempt, say, 4 transmissions to X, failing which attempt, say, 3 transmissions to A (similar to the n/m mechanism discussed in [8], and the references therein). Another possibility could be to interleave communication attempts between X and A, until any one of them is available. After, say, 7 unsuccessful attempts, the MAC layer at S may drop the packet, and notify the network layer of a route error.




Fig. 3. An example to illustrate the possibility of deafness.

Instantaneous link probing may also be applicable if nodes switch off their transceivers, to conserve power. Nodes that are known to have been awake in the recent past (either by over-hearing transmissions, or by knowing their sleeping schedules *a priori*, or through recent communication), are attempted first. A window of current history about the neighborhood channel activity is maintained in the anycast module. Neighbor selections are made after consulting this history information.

B. Power conservation

Ad hoc networks are envisioned to be networks of power constrained devices like laptops, palmtops and PDAs. Choosing appropriate next-hop neighbors, in a manner that reduces power consumption, can increase a node's lifetime. For example, the

MAC layer can choose links on the basis of its residual power level – forwarding to closer neighbors when the residual power level is small, and to further ones when residual power level is above some threshold.




Fig. 4. The scope of anycasting in conjunction with PCMA

C. Spatial Reuse

Monks *et al.* [26] have proposed a power controlled multiple access protocol (PCMA), that aims to improve spatial reuse of the channel. The key idea of PCMA is as follows. A receiver, R, informs its neighborhood about the level of additional interference that it might be able to tolerate while engaged in signal reception. A transmitter, T, in the vicinity of R, can initiate a new communication to another node, only if the interference at R due to the new communication is below R's tolerance threshold. Figure 4 illustrates this scenario. Assuming that T intends to transmit a packet to destination, D, it can choose between two possible routes, namely $\{T,N,D\}$ or $\{T,P,D\}$. Transmit power required on links T-N and P-D is large, and that on T-P and N-D is small. Observe that a routing protocol does not have any information about the interference tolerance thresholds of R, and therefore has no reason to prefer one route over the other. However, when communication between S and R is in progress, T will not be able to communicate with N – the power at which T must transmit to N can hinder signal reception at R. As a result, T must forward the packet to N only after communication between S and R has completed. This can increase the latency in routing, leading to degradation in overall network throughput.

Anycasting can be useful when using PCMA at the MAC layer. If the MAC layer at T has the option to choose between N or P as the downstream node, then T can choose to forward the packet to P. Choosing P or N as the next-hop neighbor requires knowledge of the channel conditions, available only at the MAC layer.

D. MAC-layer anycasting with directional antennas

Capacity of wireless ad hoc networks can be improved with directional antennas [27]. Several MAC and routing protocols have been proposed [12],[14],[15],[16],[7] that aim to exploit the benefits of directional beamforming. MAC-layer anycasting can be of help in enhancing the benefits even further. Consider the scenario in Figure 5. Assume that a communication between nodes E and F is in progress. To avoid interference at E, a communication between S and A may not be initiated while E is en-

gaged in communication. However, S can transmit a packet to X, without corrupting signal reception at E. The routing layer at node S remains unaware of channel status in short time scales, and is therefore incapable of making appropriate decisions on short time scales. However, if the routing layer specifies the anycast group as (A,X), S can instantaneously decide to use X as its downstream neighbor – note that S knows, to be able to transmit to A, it must wait for E to complete its transmission to F.




Fig. 5. Exploiting the benefits of beamforming using MAC-layer anycasting.

V. DESIGN TRADEOFFS

Implementing MAC-layer anycasting can introduce several tradeoffs. We discuss some of the tradeoffs in this section. We believe that the gains from anycasting will exceed its disadvantages.

A. Route optimality

Care should be taken when using MAC-layer anycasting. If the alternate next-hops specified in the anycast group correspond to routes with different “costs”, there is potential for unwanted outcomes. Figure 6 illustrates the possibility. Assume that the network layer at each intermediate node provides its MAC layer with a anycast group comprising of neighbors that are either on the fewest-hop routes, or on routes that are just one hop-count more than the fewest-hop route. Let us call this *tolerance*, which is equal to 1 hop in this example. Using a tolerance of 1 hop, for packets destined to D, the anycast group at node S will be (A,C,X) – A and C being on the fewest-hop route to D (with hop-count = 3), and X being on the next-best route to D (hop-count = 4). Similarly, the anycast group at node X will be (A,J) – A being on a 3-hop route to D while J being on a 4-hop route to D⁴. Using our proposed approach of MAC-layer anycasting, node S may forward a packet to X, which in turn may forward to J. Although these are valid local decisions at each intermediate node, observe that the overall hop-count of the traversed route will be at least five. Clearly, this exceeds the tolerance of 1 hop. Thus, without careful use, MAC-layer anycasting can cause packets to take long routes.

One possible way around this problem would be for the network layer to only specify alternate paths with identical (and minimum) costs. In the example in Figure 6, the anycast group at node S would be (A,C) – both the routes through A and C can lead to D in the fewest possible hop-counts (i.e., 3 hops). The

⁴Assume that node S is not included in the anycast group since X must not transmit a packet back to S.

anycast group at nodes A and C would then be identical – both being (B). While hop-count remains minimum, the number of alternatives in the anycast group reduces, reducing the possibility of MAC-layer anycasting.




Fig. 6. The possibility of using sub-optimal routes when using MAC-layer anycasting.

Another strategy to increase the possibility of anycasting, while restricting packet digression, can be as follows. The network layer at the source node includes the acceptable tolerance threshold within each packet. Assuming ordered anycasting, the MAC layer at each intermediate node increases a counter if it chooses to forward a packet to a neighbor that is not on one of the fewest-hop routes. Of course, to be able to increment the counter, the MAC layer needs to know whether a member of the anycast group is on the fewest-hop route. We assume that the network layer extends this information to the MAC layer by a simple grouping mechanism – the anycast group is divided into two groups, one containing neighbors on the fewest-hop routes and the other containing neighbors on longer routes. If the MAC-layer intends to forward the packet over one of the fewest-hop routes, the counter is not incremented. Clearly, at any given instant, the value of the counter (included in the packet) represents the number of additional hops that the packet has digressed. When the network layer of an intermediate node receives this packet, it must form the anycast group based on the tolerance threshold of that flow, and the current value of the counter in the packet. If the value of the counter equals the tolerance threshold, only the minimum-hop routes must be used for subsequent forwarding. Relating to the previous example in Figure 6, assume that the tolerance threshold is 1 hop. If X receives a packet (destined for D) from S, the value of the counter is already 1. Since the tolerance threshold and the value of the counter are equal, X would only choose the fewest-hop routes for subsequent forwarding – route {X,A,B,D} in this example. Clearly, digression through node J can be avoided with this particular mechanism. Observe that, if the value of the tolerance threshold is 2 hops, the network layer at X could include node J in its anycast group. Increasing the tolerance threshold reasonably, can improve performance because of the higher possibility of MAC-layer anycasting.

B. Out-of-order delivery

MAC-layer anycasting is performed on a per-packet basis. In other words, if node S intends to transmit multiple packets to D, it may choose different next-hop neighbors for forwarding each packet. Using different routes can cause packets to arrive at the destination out of order. Clearly, when using a transport

protocol like TCP, out-of-order packet delivery can be a problem [28]. Out-of-order delivery also arises when using multipath routing in any network. Other researchers are developing approaches to reduce potential degradation in TCP throughput with out-of-order delivery [28]. These approaches can be applied to MAC-layer anycasting. We intend to investigate the effects of out-of-order delivery due to MAC-layer anycasting, in our future work.

C. Source routing and MAC-layer anycasting

Several issues arise when using MAC-layer anycasting along with source routing (e.g., DSR). With source routing, the source of a packet completely specifies the route that the packet must traverse to reach its final destination. To implement MAC-layer anycasting, the source must include enough information in the header of the packets, so that intermediate nodes in the route can form their respective anycast groups, based on the available header information. A possible implementation could be to specify the alternate routes in the form of a directed acyclic graph structure, with the destination as the sink node. A node, that locates itself on the acyclic graph, can form its anycast group by selecting its downstream neighbors from the graph. Of course, the header length can increase significantly, adding to the control overhead associated to routing. Discovering multiple routes can also increase control overhead. In “basic” DSR, a node drops all duplicate route request (RREQ) packets that it receives during the route discovery phase. To facilitate discovery of multiple routes, a node may need to forward one or more duplicate RREQ packets. In addition, the destination node must reply with multiple route reply (RREP) packets, each carrying a distinct route to the source. The net gains due to MAC-layer anycasting, offset by the disadvantages of increased control overhead, is a tradeoff we plan to evaluate.

VI. CONCLUSION

We propose MAC-layer anycasting for ad hoc wireless networks. The network layer specifies multiple downstream nodes, from which the MAC layer chooses a suitable node, based on instantaneous network conditions. We believe that MAC-layer anycasting can be combined with different routing schemes – the routing protocol can form the anycast group based on its metric of choice. Anycasting can then be applied to this anycast group. We illustrate cases in which anycasting may offer performance gain. We discussed some specifics of implementation and discussed the performance tradeoffs that arise due to MAC-layer anycasting. Evaluating the performance of anycasting through simulations is a topic for future work.

REFERENCES

- [1] J. Broch, D. Maltz, D. Johnson, Y. Hu, and J. Jetcheva, “A performance comparison of multi-hop wireless ad hoc network routing protocols,” in *Mobile Computing and Networking*, 1998, pp. 85–97.
- [2] E. Royer and C. Toh, “A review of current routing protocols for ad-hoc mobile wireless networks,” 1999.
- [3] D. Johnson and D. Maltz, “Dynamic source routing in ad hoc wireless networks,” in *Mobile Computing*, Kluwer Academic Publishers, 1996.
- [4] Charles E. Perkins and Pravin R. Bhagwat, “Highly dynamic destination-sequenced distance vector routing (DSDV) for mobile computers,” in *Proceedings of ACM SIGCOMM*, 1994, pp. 234–244.
- [5] G. Pei, M. Gerla, and T. Chen, “Fisheye State Routing in Mobile Ad Hoc Networks,” in *ICDCS Workshop on Wireless Networks and Mobile Computing*, Taipei, Taiwan, April 10 2000.
- [6] S. Roy, D. Saha, S. Bandyopadhyay, T. Ueda, and S. Tanaka, “A network-aware MAC and routing protocol for effective load balancing in ad hoc wireless networks with directional antenna,” in *Proceedings of ACM MobiHoc*, June 2003, pp. 88–89.
- [7] S. Bandyopadhyay, K. Hausike, S. Horisawa, and S. Tawara, “An adaptive MAC and directional routing protocol for ad hoc wireless networks using ESPAR antenna,” in *Proceedings of ACM MobiHoc*, October 2001, pp. 243–246.
- [8] M. Pursley and H. Russell, “Network protocols for frequency-hop packet radios with decoder side information,” *IEEE Journal on Selected Areas in Communications*, vol. 12, no. 4, pp. 612–621, May 1994.
- [9] P. Larsson, “Selection diversity forwarding in a multihop packet radio network with fading channel and capture,” *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 5, no. 4, pp. 47–54, 2001.
- [10] S. Jain, Y. Lv, and S. Das, “Exploiting path diversity in link layer in wireless ad hoc networks,” in *manuscript (personal communication)*, July 2003.
- [11] IEEE 802.11, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, 1999.
- [12] Y. Ko, V. Shankarkumar, and N. H. Vaidya, “Medium access control protocols using directional antennas in ad hoc networks,” in *Proceedings of IEEE INFOCOM*, March 2000, vol. 1.
- [13] A. Nasipuri, S. Ye, J. You, and R. Hiroto, “A MAC protocol for mobile ad hoc networks using directional antennas,” in *Proceedings of IEEE Wireless Communication and Networking Conference*, September 2000, vol. 3, pp. 1214–1219.
- [14] R. Ramanathan, “On the performance of ad hoc networks with beam-forming antennas,” in *Proceedings of ACM MobiHoc*, October 2001, pp. 95–105.
- [15] M. Takai, J. Martin, R. Bagrodia, and A. Ren, “Directional virtual carrier sensing for directional antennas in mobile ad hoc networks,” in *Proceedings of ACM MobiHoc*, June 2002, pp. 39–46.
- [16] R. Roy Choudhury, X. Yang, N. H. Vaidya, and R. Ramanathan, “Using directional antennas for medium access control in ad hoc networks,” in *Proceedings of ACM MOBICOM*, September 2002.
- [17] T. Korakis, G. Jakllari, and L. Tassioulas, “A MAC protocol for full exploitation of directional antennas in ad hoc wireless networks,” in *Proceedings of Mobihoc*, 2003.
- [18] X. Qin and R. Berry, “Exploiting multiuser diversity for medium access control in wireless networks,” in *Proceedings of INFOCOM*, 2003.
- [19] M. Tsatsanis and R. Zhang, “Network assisted diversity for random access wireless networks,” *IEEE Transactions on Signal Processing*, vol. 48, no. 3, March 2000.
- [20] D. Decouto, D. Aguayo, J. Bicket, and R. Morris, “A high-throughput path metric for multi-hop wireless networks,” in *Proceedings of MobiCom*, 2003.
- [21] M. Yarvis, W. Conner, L. Krishnamurthy, J. Chhabra, and B. Elliott, “Real-world experiences with an interactive ad hoc sensor network,” in *Proceedings of IEEE ICPPW*, 2002.
- [22] S. Xu and T. Saadawi, “Does the IEEE 802.11 MAC protocol work well in multihop wireless ad hoc networks,” *IEEE Communications Magazine*, June 2001.
- [23] R. Roy Choudhury and N. H. Vaidya, “Deafness: A MAC problem in ad hoc networks when using directional antennas,” Tech. Rep., Coordinated Science Laboratory, and University of Illinois at Urbana-Champaign, July 2003.
- [24] R. Zheng and J. Hou, “Asynchronous wakeup for ad hoc networks,” in *Proceedings of MobiHoc*, June 2003.
- [25] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris, “SPAN: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks,” in *Mobile Computing and Networking*, 2001, pp. 85–96.
- [26] Jeffery P. Monks, Vaduvur Bhargavan, and Wen-Mei W.Hwu, “A power controlled multiple access protocol for wireless packet networks,” in *Proceedings of INFOCOM*, 2001, pp. 219–228.
- [27] S. Yi, Y. Pei, and S. Kalyanaraman, “On the capacity improvement of ad hoc wireless networks using directional antennas,” in *Proceedings of Mobihoc*, 2003.
- [28] E. Blanton and M. Allman, “On making TCP more robust to packet re-ordering,” *ACM Computer Communication Review*, Jan. 2002.