

MASTER COPY: PLEASE KEEP THIS "MEMORANDUM OF TRANSMITTAL" BLANK FOR REPRODUCTION PURPOSES. WHEN REPORTS ARE GENERATED UNDER THE ARO SPONSORSHIP, FORWARD A COMPLETED COPY OF THIS FORM WITH EACH REPORT SHIPMENT TO THE ARO. THIS WILL ASSURE PROPER IDENTIFICATION. NOT TO BE USED FOR INTERIM PROGRESS REPORTS; SEE PAGE 2 FOR INTERIM PROGRESS REPORT INSTRUCTIONS.

MEMORANDUM OF TRANSMITTAL

U.S. Army Research Office
ATTN: AMSRL-RO-BI (TR)
P.O. Box 12211
Research Triangle Park, NC 27709-2211

- | | |
|---|--|
| <input checked="" type="checkbox"/> Reprint (Orig + 2 copies) | <input type="checkbox"/> Technical Report (Orig + 2 copies) |
| <input type="checkbox"/> Manuscript (1 copy) | <input type="checkbox"/> Final Progress Report (Orig + 2 copies) |
| | <input type="checkbox"/> Related Materials, Abstracts, Theses (1 copy) |

CONTRACT/GRANT NUMBER: **W911NF0410224 (46637-CI-MUR)**

REPORT TITLE: Paolo Casari, Marco Levorato, Daniele Mazzi, Michele Zorzi, "On the Design of Routing Protocols for MIMO Ad Hoc Networks Under Uniform and Correlated Traffic", accepted at the International Wireless Communications and Mobile Computing Conference (IWCMC), Crete, Greece, August 6-8, 2008.

is forwarded for your information.

SUBMITTED FOR PUBLICATION TO (applicable only if report is manuscript):

Sincerely,

Michele Zorzi
University of California, San Diego

REPORT DOCUMENTATION PAGE

Form Approved
OMB NO. 0704-0188

Public Reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comment regarding this burden estimates or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188,) Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE July 2008	3. REPORT TYPE AND DATES COVERED Reprint 01 June 2004 - 31 March 2007	
4. TITLE AND SUBTITLE On the Design of Routing Protocols for MIMO Ad Hoc Networks Under Uniform and Correlated Traffic			5. FUNDING NUMBERS W911NF0410224	
6. AUTHOR(S) Paolo Casari, Marco Levorato, Daniele Mazzi, Michele Zorzi			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of California. San Diego 9500 Gilman Drive, La Jolla, CA 92093-0407			10. SPONSORING / MONITORING AGENCY REPORT NUMBER N/A	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U. S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			11. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.	
12 a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.		12 b. DISTRIBUTION CODE N/A		
13. ABSTRACT (Maximum 200 words) In this paper we investigate the effects of different routing choices in a MIMO ad hoc network. The MIMO physical layer provides a number of degrees of freedom that a correct protocol design could exploit to the network's advantage. To illustrate this point, we compare two approaches. The first is fixed shortest path forwarding, with routes computed a priori. The second involves a dynamic relay choice based on a specific metric. This metric has been designed to yield an overall assessment of the relay's ability to forward packets. A key point of our design lays in its cross-layer nature, that allows to harvest the advantages coming from the interaction of the physical layer and the upper layers. To this end, we propose here a technique to track the receive capabilities of eligible forwarders so as to avoid that they become overloaded and unable to demultiplex superimposed waveforms. We carry out a detailed performance evaluation to assess the improvements provided by our solution with respect to the baseline shortest path routing protocol, and conclude that a standard design paradigm, that splits routing and lower-level protocols, severely limits performance.				
14. SUBJECT TERMS			15. NUMBER OF PAGES 6	16. PRICE CODE N/A
17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION ON THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

NSN 7540-01-280-5500

Standard Form 298 (Rev.2-89)
Prescribed by ANSI Std. Z39-18
298-102

On the Design of Routing Protocols for MIMO Ad Hoc Networks under Uniform and Correlated Traffic

Paolo Casari*, Marco Levorato*, Daniele Mazzi* and Michele Zorzi*[‡]

*Department of Information Engineering, University of Padova, Italy. E-mail: {firstname.lastname}@dei.unipd.it

[‡]California Institute of Telecommunications and Information Technology, UC San Diego, USA

Abstract—In this paper we investigate the effects of different routing choices in a MIMO ad hoc network. The MIMO physical layer provides a number of degrees of freedom that a correct protocol design could exploit to the network’s advantage. To illustrate this point, we compare two approaches. The first is fixed shortest path forwarding, with routes computed a priori. The second involves a dynamic relay choice based on a specific metric. This metric has been designed to yield an overall assessment of the relay’s ability to forward packets. A key point of our design lays in its cross-layer nature, that allows to harvest the advantages coming from the interaction of the physical layer and the upper layers. To this end, we propose here a technique to track the receive capabilities of eligible forwarders so as to avoid that they become overloaded and unable to demultiplex superimposed waveforms. We carry out a detailed performance evaluation to assess the improvements provided by our solution with respect to the baseline shortest path routing protocol, and conclude that a standard design paradigm, that splits routing and lower-level protocols, severely limits performance.

Index Terms—MIMO ad hoc networks, routing, dynamic relay selection, cross-layer design.

I. INTRODUCTION

Ad hoc networks are collections of radio terminals that perform packet-radio communications through a shared radio channel, without the mediation of any infrastructure. Even though completely ad hoc networks are not yet widespread, infrastructured wireless networks (as epitomized by the IEEE 802.11 protocol [1]) are now part of many people’s everyday life. Thus, how to improve radio access proficiency and to make available a larger communication bandwidth and bit rate are currently hot research topics. Recently, multipacket reception (MPR) in ad hoc networks has emerged as an important enabling technology to improve the communication efficiency by reducing congestions and traffic bottlenecks in wireless networks. MPR allows the receptions of multiple packets in parallel over the same radio channel. There are a number of techniques that can be employed to superimpose multiple packets. For example, different CDMA waveforms could be assigned to different users, having the receiver separate them [2], [3]. Alternatively, or in addition, signal processing techniques that involve the use of antenna arrays could be applied to demultiplex superimposed signals [4]. In this second case, the system is also defined as Multiple Input–Multiple Output (MIMO). While MIMO systems are not necessarily coupled with CDMA, the use of both may prove to be very helpful, depending on incoming traffic conditions. For example, consider a V-BLAST system [5], where multiuser detection is performed through successive interference cancellation [6]. The use of CDMA yields additional degrees of freedom, allowing to demultiplex a number of incoming signals that would not be separable otherwise using MIMO or CDMA alone.

In this paper, we focus on a MIMO–CDMA multihop ad hoc network. We deploy a cross-layer design that involves the

Medium Access Control (MAC) and physical (PHY) layer, and compare different routing solutions over such a network. Our primary aim is to show how routing impacts this kind of network, and to show how a correct cross-layer design can manage link activation and data reception according to both PHY requirements and routing needs, ultimately improving performance substantially. In fact, the design of routing in MIMO networks is different with respect to standard ad hoc networks. MIMO yields potentially large communications benefits, thanks to spatial diversity, spatial multiplexing, and interference cancellation at the receiver. However, these benefits require proper protocols to be fully exploited. For example, a number of tradeoffs arise, *e.g.*, between the amount of transmit antennas in the neighborhood of a receiver, and the protection that receiver will achieve against interfering transmissions [7]. Similarly, when employing more transmit antennas under a global transmit power constraint, the expected reach of the transmission decreases, thus increasing the number of hops that would be required on a multihop path. In this paper, we address these tradeoffs through the design of a proper routing protocol that takes more advantage of MIMO communications. Our solution is compared to a baseline shortest path forwarding strategy by means of extensive simulation results, in order to assess the improvements granted by our specific protocol design.

II. RELATED WORK

Protocol design for MIMO ad hoc networks constitutes an emerging topic that is currently being investigated by a number of research groups. Since the effort is still in its infancy, the main research direction stems from the work on the characterization of MIMO links, and tries to apply this knowledge to ad hoc networks in an efficient manner [8]. For the same reason, there is currently little work on routing in MIMO ad hoc networks. However, some contributions have provided first insights into the topic.

A graph theory approach is used in [9] to find a transmission schedule in the network that maximizes the utility of MIMO features such as spatial multiplexing and interference suppression. In [10], the authors investigate optimum link schedules that carry flows in a multihop network at the requested rate, under throughput feasibility and minimum power constraints. The work in [11] is also based on joint power control and routing, and focuses on the allocation of channel bandwidth to each node on a multihop path, when using orthogonal channels for each active link. Along the same line, [12] explores a joint routing and link scheduling optimization to find the optimal tradeoff between energy and latency. A clustered MIMO ad hoc network is considered where nodes in the same cluster cooperate in communications. The work in [13] presents a simulation study of the tradeoff between spatial multiplexing and spatial reuse for varying number of antennas. The study considers the joint effect of multiple network parameters such as hop length, node density and transmission range, yet assumes channel knowledge

at the transmitter, which may be unrealistic in some scenarios. An interesting work on MAC and routing in MIMO networks is found in [14]. There, the authors extend an 802.11-like protocol to exploit spatial diversity through space-time codes [15], and analyze the impact of MAC on routing by evaluating the relation between incurred delay and per-hop advancement.

These papers are interesting in that they assess many problems and challenges related to the design of MIMO ad hoc networks. Still, there are no works to date that present a full solution for cross-layer networking in this scenario. Our purpose is to draw some conclusions on the way routing protocols affect MIMO networks, and how they should be designed to yield acceptable performance.

III. SYSTEM DESCRIPTION

In order to maintain the paper self-contained, in this Section we provide a brief overview of the detection algorithm at the physical layer. As the main focus of our work is on the interactions between MAC and routing protocols, these interactions will be discussed to a broader extent. The interested reader is referred to [16], [17] for a more thorough discussion on PHY details.

In our network, each terminal is equipped with multiple antennas. During a transmission phase, each node divides the packets in packet data units (PDUs) of fixed length. Each node has a unique spreading sequence that is used to construct CDMA waveforms for all PDUs. Each CDMA PDU is then simultaneously transmitted through a different antenna, one PDU per antenna, using the well known V-BLAST algorithm [5]. In general, a receiver has to demultiplex a number of superimposed PDUs, possibly coming from different transmitters. We assume that each PDU contains a training sequence that the receiver can use to estimate channel effects as the signals propagate from the transmit antennas to each of its receive antennas. Furthermore, we assume that channel effects are constant through the whole transmission of the PDUs. The standard V-BLAST algorithm is adapted to exploit CDMA signals as in [6]. Briefly, after applying a space-code-matched filter to the received signals, the detection process consists of subsequent steps. At each step, the most powerful signal still to detect is selected, and proper weights are chosen that allow to isolate that signal among all received signals according to some criterion, such as zero forcing. Then, the signal is detected and canceled from the overall received signal. This process continues by selecting progressively weaker signals until all are detected.

A PHY such as the one described before has MPR capabilities that can be proficiently exploited, in order to receive wanted packets while still protecting them from strongest interference (that should be detected and canceled, instead). As a node has limited channel estimation resources, it is very important to choose a suitable subset of signals to detect and/or cancel. The MAC protocol components have the twofold purpose to favor this choice while still enforcing simultaneous channel access, so as to exploit the MPR at the receiver. To this end, we assume that all nodes communicate over subsequent frames, and that they are frame-synchronous. Each frame is divided into four slots, a Request-To-Send (RTS), a Clear-To-Send (CTS), a data, and an acknowledgment (ACK) slot. Unlike in 802.11, RTSs and CTSs are not used to set up links while blocking neighboring transmissions, but instead to pass information about transmission requests to the neighbors. While the data and ACK phases are straightforward (all PDUs allowed by CTSs are transmitted and their correct reception confirmed, respectively), the RTS and CTS phases require a specific policy to properly drive channel access.

A. RTS policy

The RTS policy is required to enforce transmission parallelism without overloading a link with too many simultaneous signals. To this end, each neighbor is assigned a “class,” namely the maximum number of antennas that can be used when transmitting to a set of nodes including that neighbor, in order to guarantee some minimum prescribed performance. The choice of the class can be based on various considerations involving, *e.g.*, the distance of the node and the history of the probability of error over that link. We will elaborate more on this point later: for now we assume that the class of each neighbor is known. When forming an RTS, a transmitter scans its backlog, to read the destination and the number of PDUs left to transmit for that packet. Iteratively, the transmitter checks if transmitting the required PDUs violates any constraint posed by the class any other node addressed by the RTS. If at least a subset of these PDUs can be transmitted, the sender writes a request to the destination node for the specific number of PDUs in the RTS message; otherwise it leaves the packet in the queue for a later attempt. The RTS composition ends when adding further requests would exceed a class limit. For example, if the RTS currently contains a request for 4 PDUs toward a class 8 neighbor, the transmitter is allowed one more simultaneous request for up to 4 PDUs toward another class 8 neighbor. Conversely, in the same conditions, it cannot transmit any PDU toward a class 2 neighbor, because the number of PDUs requested to the first one (4) exceeds the class of the latter (2).

B. CTS policy

Typically, MPR capabilities at the receiver allow the correct detection of a number of RTSs, because these packets are small and transmitted at full power using one antenna. This endows the receiver with some knowledge of neighboring transmissions, so that it can choose what to grant or block (among the requests directed to itself) and what to defend against (among interfering transmissions that could take place). To make a proper decision, the node must also consider that *i*) it has limited channel estimation capabilities (for PDU detection and cancellation), and *ii*) that it cannot block transmissions directed toward somebody else. The way these decisions are made is a CTS policy. A convenient way to grant transmissions is the “follow traffic” (FT) policy [17], which requires the node to grant at least one request directed toward itself, and to consider the remaining requests in order of decreasing received power. Basically, if the request is meant for the node, it grants it in the CTS, otherwise the node pre-allocates channel estimation resources needed to detect and cancel the signal, as it will be a potentially powerful interferer. This process goes on until there are either no more requests, or no more channels can be estimated, and the CTS is sent thereafter.

After all receivers have sent their CTSs, the transmitters send data altogether according to the information contained in the CTSs, and wait for an ACK by the receiver to confirm a correct data reception on a per PDU basis. In case no CTS is received, or some of the addressed destinations have not sent a CTS, a backoff policy is applied. This is necessary to avoid useless persistence in channel access, and allows a possible local congestion near the receiver to be solved before trying again. We consider two different backoff policies. The first one, Node-Lock, is more conservative as it blocks all transmissions toward any other node. The second one, Dest-Lock, blocks transmissions toward the only neighbor that did not send the CTS, and is thus more aggressive. In either case, the duration of the backoff is a number of frames uniformly distributed in the interval $[1, B_{max}]$, where B_{max} increases exponentially in the number of subsequent failures.

IV. ROUTING

Before proceeding with considerations on routing performance on this MIMO ad hoc network, let us sum up the main features that affect (or are related to) routing. Firstly, all protocols encourage spatial multiplexing, to exploit MPR capabilities at the physical layer. There are three kinds of limits to transmissions: the class of the neighbors, which bounds the maximum spatial parallelism during a data exchange; the CTS policy, which controls traffic based on throughput, power and interference; the backoff policy, which defers transmission directed to some or all neighbors. These control mechanisms are indeed necessary [7], [8] to ensure link reliability, yet they may interfere with commonly used routing algorithms if they are straightforwardly applied to a MIMO system.

Routing needs are different than in single-user detection wireless networks. For instance consider a local congestion in a certain part of the network. The FT algorithm would control traffic and allow data to get through even in the presence of strong interference, but at the price of an actual shortening of the transmission range. Recall that FT basically drives signal detection in order of decreasing received power. This means that pursuing shortest path routing, whereby a packet travels a long distance on a single hop, may not be the best policy. In fact, long distance signals bear a lower power and might suffer from less frequent transmission opportunities and longer delays due to the way FT operates.

While this is just an example of a critical situation, it conveys the main concept behind a correct routing algorithm, *i.e.*, the need to make routing decisions based on parameters measured at the PHY and MAC layers (and/or their history). Suppose that a node knows its neighbors. In our MIMO ad hoc network, MPR of RTSs over the initial part of a frame allows to understand which neighbors are transmitting and which are free. Then a node can choose, among the free ones, the neighbor that would offer the best “utility” if relaying the message. To this end, it is a good choice to mix some relevant metrics such as link reliability, actual advancement, capability to serve the packet (*e.g.*, long or short queue) and so forth. We choose a mix of four utility figures. Let \mathcal{N} be the set of neighbors of the transmitter that are free during the current frame. Call $i \in \mathcal{N}$ one of the neighbors and let $\xi_i = \frac{d_{S,D} - d_{i,D}}{d_{S,D}}$ be the advancement provided by i , *i.e.*, the distance reduction that i would provide, normalized to the distance between the source S and the destination D . Let also c_i be the class of node i (the maximum number of antennas allowed when transmitting to a set of neighbors that includes i). Finally, let q_i and b_i be i 's queue level and residual backoff timer, respectively. While it is assumed that the nodes know their own distance and the destination's, all other information is kept updated by reading the content of correctly received signaling messages. The overall utility yielded by the choice of a neighbor i is then computed as

$$U(i) = \alpha_1 f_\xi(\xi_i) + \alpha_2 f_c(c_i) + \alpha_3 f_q(q_i) + \alpha_4 f_b(b_i), \quad (1)$$

where the functions $f_\xi(\cdot)$, $f_c(\cdot)$, $f_q(\cdot)$, and $f_b(\cdot)$ can be tuned to properly account for each parameter and are assumed to return a value in the interval $[0, 1]$, where 1 is best. A specific choice for these functions will be given in section V-A. All functions are accounted for in $U(i)$ through a convex combination with coefficients α_j , *i.e.* $0 \leq \alpha_j \leq 1$ and $\sum \alpha_j = 1$. It is important to jointly account for all these metrics in order to conjure up a complete picture of the relay's status at a glance. A correct balance between the four weights allows to choose a relay that has desirable features when routing in our MIMO network. For example, choosing only based on advancement tends to create

shortest paths, but the receivers will have to handle low power signals that the FT policy will most likely block. On the other hand, routing only based on the queue or the backoff timer of the relay, without considering advancement, would let the packet enter loops. It is also important to account for the capacity of the link, a measure of which is provided by its class. In fact, the class is dynamically updated based on the relaying performance of the neighbor, and higher classes tend to represent less loaded and more reliable nodes. All of these considerations should be kept in mind when designing the class update policy and the function $f_c(c_i)$, whose details for our setting are given in Section IV-A. Finally, the queue level and the residual backoff timer are a measure of the “ability” of the relay to further forward the packet (the longer the queue and backoff, the longer the delay incurred).

A. Class update scheme

The definition of a node's class is one of the elements that tightly binds our MAC and PHY layers, and therefore deserves deeper attention. Recall that by class of a node we mean the maximum number of antennas allowed when transmitting to a set of neighbors which includes that node. The reliability and efficiency of a transmission critically depend on the class. Overloading a link that is incapable of supporting high spatial multiplexing generates errors, and underutilizing a potentially good link limits throughput. There is also an important interaction with FT. If many transmissions take place on high-class links, generally the number of requests will be high, thus FT will try to protect wanted data from interference by granting transmissions to only a few senders. If everybody does that, only a few CTSs are sent, and throughput is limited as well. The main message, here, is hence that we need to balance the link classes correctly and keep them updated. A good way to do this is to allow the class to take values in a limited set, *e.g.*, $\mathcal{C} = \{2, 4, 6, 8\}$, where class 1 is not chosen in order to enforce spatial multiplexing. Given a neighbor i , the class value is derived by tracking the success ratio (or reliability) over the link toward i , according to the first-order recursive formula

$$r'_i = a n_i^{(a)} / c_i + (1 - a) r_i, \quad (2)$$

where n_i is the number of correctly received PDUs among those sent to i , and c_i is the current i 's class. This value is then compared to two thresholds T_d and T_u . If $r'_i < T_d$ the class is decreased to the next smaller value in \mathcal{C} . Conversely, if r'_i exceeds T_u , the class is increased to the next greater value. Note that, calling $n_i^{(s)}$ the number of PDUs sent to i , the $n_i^{(a)} / c_i$ can also be written as $n_i^{(a)} / n_i^{(s)} \cdot n_i^{(s)} / c_i$, *i.e.*, as the success ratio of the transmission to i times the ratio of the sent PDUs to the node class (*i.e.*, the actual “use” of the class). This way, the class of a node is likely to be reduced even if the class itself is underused due to a limited number of grants by FT. In turn, this prevents the FT drawback detailed before, whereby too many nodes require the transmission of too many PDUs and FT allows just a few grants to prevent interference, ending up limiting throughput as well.

V. RESULTS

A. Simulation setting

In order to evaluate the performance of the dynamic routing solution outlined in section IV, we arrange 64 nodes in an 8×8 square grid, where the distance between nearest neighbors is 25 m. Each node is equipped with an array of 8 antennas. All antennas are used during reception, whereas transmissions are performed with as many antennas as needed to send the PDUs




Figure 1. One-hop throughput vs. traffic for static shortest path routing, for both Dest-Lock and Node-Lock.




Figure 2. End-to-end throughput vs. traffic for static shortest path routing, for both Dest-Lock and Node-Lock.

granted by the received CTSs. Unless otherwise specified, traffic is generated according to a Poisson process with parameter λ packets per second per node. All nodes can generate traffic. Each packet is composed of k 1000-bit PDUs, with k uniformly chosen in the set $\{1, 2, 3, 4\}$. For each generated packet, the destination is randomly chosen to be one of the other nodes in the network. The packets waiting to be transmitted are temporarily stored in a queue that can hold up to 512 1000-bit PDUs. Furthermore, one half of the queue is reserved for packets coming from other nodes to be relayed, so that there is always a certain amount of resources devoted to the routing process. The other half is general-purpose, *i.e.* it can be used for both node-originated packets and for packets to be relayed. Packets that cannot be successfully transmitted are dropped from the queue after a timeout interval.

All communications take place according to the frame structure described in Section III. The length of all signaling packets is fixed to 200 bits. During the RTS, CTS and ACK phases, the signaling packets are sent using one antenna (each node sends one cumulative packet containing all requests/grants/acknowledgments). This minimizes the number of packets sent per node while keeping the transmit power high, so that there is a high probability that all signaling messages are correctly received. Moreover, recall that the data phase involves the parallel transmission of multiple PDUs, each sent through one antenna. The length of a PDU is fixed to 1000 bits, hence the data phase length is fixed as well.

The channel estimation process required to correctly detect spatially multiplexed signals is carried out using a limited amount of estimation resources. More specifically, each node can estimate at most 32 different training sequences. That means that it can (try to) detect at most 32 signaling packets coming from the closest 32 neighbors, or 32 PDUs during the data transmission phase. As described before, the FT policy is in charge of choosing which PDUs are worth being detected among the requests presented in the signaling messages.

The different routing approaches that we compare in the following involve a shortest-path routing algorithm and a dynamic relay selection algorithm. The shortest path routing is meant as a baseline for the comparison. In this case, the paths are pre-calculated using the Dijkstra algorithm, with physical distance as the link weight. For the dynamic selection, we apply the considerations in section IV. More specifically, we set the normalizing functions in (1) as follows. Given a node $i \in \mathcal{N}$, where \mathcal{N} is the set of neighbors of the current sender, $f_{\xi}(\xi_i) = \xi_i / \max_{j \in \mathcal{N}} \xi_j$, $f_c(c_i) = c_i/8$ (8 is the maximum node class), $f_q(q_i) = \rho \cos^2(2\pi q_{R_i}/256) + (1 - \rho) \cos^2(2\pi q_{N_i}/256)$

(where q_{R_i} and q_{N_i} are the level of the reserved and general-purpose queue, respectively, and $\rho = 0.75$ to favor the routing process¹). Note that we have warped the information on the queue length through a square cosine function. The reason is that relays with a low number of packets in queue are typically as able to forward packets as well as nodes with an almost empty queue. Hence we wish to give a higher penalty to full queues (that cause longer delays) while considering empty or low queues almost equivalent. Finally $f_b(b_i) = 1 - b_i / \max_{j \in \mathcal{N}} b_j$, and the thresholds for increasing and decreasing the class have been set to $T_u = 0.75$ and $T_d = 0.25$, respectively. The weighing scalars for the four functions have been set to $\{\alpha_1, \alpha_2, \alpha_3, \alpha_4\} = \{0.4, 0.2, 0.25, 0.15\}$. The values aim at balancing the involved factors according to the observations in section IV, and have been derived through extensive simulations.

B. Static shortest-path routing

Let us now consider the static routing scenario where all paths are pre-computed, and each hop travels the farthest possible distance. The first result we present in Figures 1 and 2 is an evaluation of local (*i.e.*, 1-hop) and end-to-end throughput for the two different backoff techniques discussed in Section III. Our purpose here is to show the benefit yielded by the automatic update of the classes with respect to a system that blindly assumes that all neighbors have class 8. First, note that the introduction of a dynamic class update is definitely beneficial for both local and end-to-end throughput. No path is being updated here (recall routing paths are pre-computed). Simply, each transmitter tracks the changes in the reliability of the links according to (2), and updates the class.

The first thing to note is the difference among the two backoff policies. Dest-Lock reaches a higher throughput than Node-Lock, both local and end-to-end. Moreover, Node-Lock's throughput decays when congestion builds up, whereas Dest-Lock's tends to remain stable. The reason behind this is that Node-Lock refrains from all transmissions at the first error that takes place. While a node cannot transmit, it can still receive and generate packets, thereby increasing its queue level. If subsequent backoffs take place, the node gets rapidly congested, and in particular all the paths that get through it undergo a great performance decrease. The static routing protocol considered here does not help at all to relieve this situation. This throughput performance is counterintuitive, as one would expect that an aggressive transmission

¹This value has been chosen through extensive simulations to yield good relay choices in our scenarios.

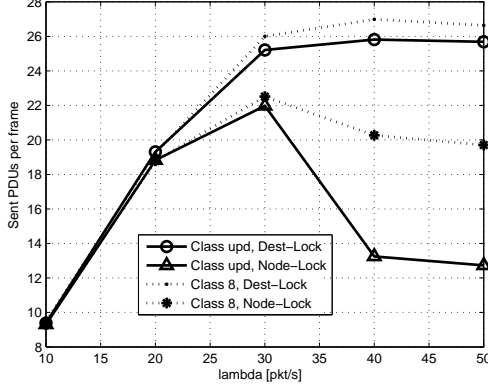


Figure 3. Average number of PDUs sent per frame, for both Dest-Lock and Node-Lock.

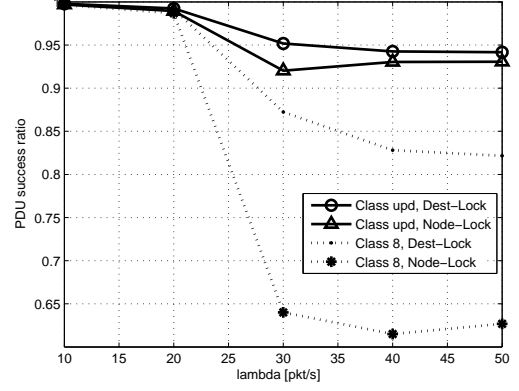


Figure 4. Average success ratio for a PDU transmission, for both Dest-Lock and Node-Lock.

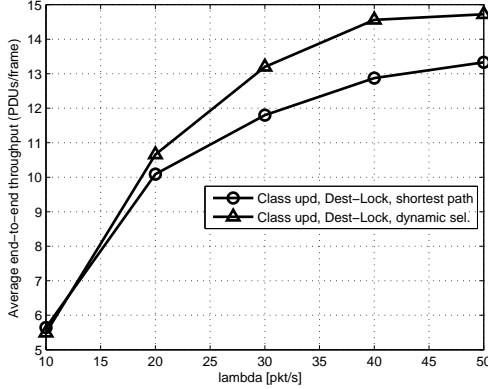


Figure 5. End-to-end throughput with dynamic and static relay selection under correlated traffic.




Figure 6. End-to-end throughput with dynamic and static relay selection under uniform traffic.

policy performs worse. Let us explain this result with the help of Figures 3 and 4, that represent the average number of PDUs sent per frame and the average success ratio, respectively. From Figure 3 we can see that Dest-Lock, on average, sends more PDUs per frame than Node-Lock, which is expected. Also, assuming that all neighbors have class 8 also increases the number of transmissions. Figure 4 suggests that the class update is beneficial from the point of view of success ratio. Regardless of the backoff policy, a dynamic class update keeps the success ratio consistently above 90%. Otherwise, it would drop, and even more dramatically for Node-Lock. The counterintuitive result, here, is that Node-Lock experiences a lower success ratio, even though fewer PDUs per frame are sent. To explain this fact, recall that with Node-Lock, the nodes exiting backoff can transmit to any neighbor. Since they keep receiving while in backoff, they will have more traffic to handle when they can transmit again, and are thus more likely to inject more requests at once in the network. Especially at high traffic, the nodes tend to back off more often, therefore the overall number of PDUs sent is kept small, as seen from Figure 3, and many transmissions are granted. However, due to shortest path routing, many of these transmissions are performed towards far nodes, and thus bear a low power and are more difficult to detect. This is confirmed by the success ratio trends in Figure 4. The situation gets worse as traffic increases, because the nodes stay more in backoff, which keeps the throughput low. Since this does not improve the link reliability, the class update policy tends to decrease the class of the neighbors. In turn, this significantly improves the reliability (roughly from 62% to 93% at high traffic), but limits transmissions and keeps the throughput very low, even

slightly worse than with the non-updating policy.

C. Dynamic relay selection and correlated traffic generation

The previous study suggested that class updating is a convenient transmission control element, and that the Dest-Lock policy performs better than Node-Lock in a routing scenario. Hence, we will consider only the former in the forthcoming study. The purpose of this section is to show the impact of a different forwarding scheme where the relay is dynamically chosen. In particular, we wish to carry out this comparison in a more general scenario, where the traffic is not uniform. This case is more realistic as traffic is usually generated on-demand, for example when new data sessions are opened. Furthermore, it is a better test for our adaptive routing protocol and class updating policy, that should help redistribute high volumes of traffic when they are injected in the network.

We considered the correlated traffic model presented in [18], that is described as follows. Each node is assumed to have a number of open data sessions. Each session generates a number of flows, each being formed by a certain number of packets. The model assumes that the sessions are created according to a Poisson process with rate λ_s . Sessions begin at the time they are created, and finish when the last packet of the last flow has been generated. The number of flows belonging to a session is assumed to be a random variable with geometric distribution and average N_f . All flows are generated as time elapses in accordance to another Poisson process with rate λ_f . Similarly, each flow originates a random number of packets, with geometric distribution and average N_p , which are generated according to a

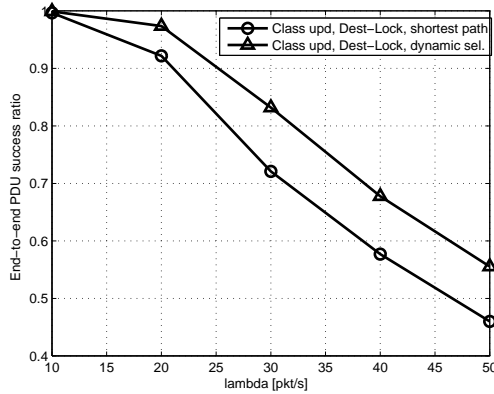


Figure 7. Average end-to-end delivery ratio for one PDU under correlated traffic.

Poisson process with rate λ_p . The average packet generation rate is thus $\lambda = \lambda_s N_f N_p$. Whenever a node starts a new session, a recipient is randomly chosen among the other network nodes. This recipient will be sent all packets originated within the session. For our study we have chosen $\lambda_f = 2$, $N_f = 2$, $\lambda_p = 20$, $N_p = 5$. This makes sessions bursty, requiring the nodes to make full use of their forwarding capability in order to send data efficiently.

The first results are shown in Figures 5 and 6, that depict the average end-to-end throughput comparing static (shortest path) and dynamic relay selection, under correlated and uniform traffic, respectively. The selection of a different relay yields a substantial benefit in the correlated traffic scenario. In fact, choosing a different relay each time ensures that the packets can be directed toward more convenient network regions, where the nodes have a higher probability to forward them further toward the destination. The adaptivity allows to process even bursty traffic volumes. The advantage is less visible under uniform traffic, where all nodes are almost equally loaded, and the metric for relay selection is likely to perceive the same utility for all neighbors. The use of dynamic path selection leads to better network connectivity, because more packets are preserved as they travel longer routes. Figures 7 and 8 also support this claim by showing the average end-to-end success ratio and the number of PDUs received after more than one hop respectively. In this case we consider only correlated traffic. In both cases, the use of dynamic relay selection allows more data to get through. For example, 10% more packets can correctly reach the destination, as the number of PDUs that travel more than one hop at high traffic is increased of 1.5.

VI. CONCLUSIONS

In this paper, we have discussed the issues and tradeoffs that arise when designing routing protocols for MIMO ad hoc networks. We have compared a dynamic relay selection approach to a baseline shortest path strategy. From the presented results, we can conclude that an adaptive policy that chooses relays dynamically and tunes their class (*i.e.*, allowed level of spatial multiplexing) yields better results in our MIMO ad hoc network. In fact, the complexity of such a setting can be more profitably controlled through a cross-layer design of the various components. For this reason, a straightforward use of fixed shortest path routes has proved to be a working, but ineffective solution. The adaptive forwarding policies offer better performance because they can harvest network resources more efficiently by conveniently redistributing traffic over the nodes.

Future work on this topic includes an evaluation of routing performance in more complex scenarios as well as modeling the




Figure 8. Average number of PDUs received after more than one hop under correlated traffic.

performance of MIMO routing protocols under imperfect channel estimation.

REFERENCES

- [1] IEEE Standards Department, *ANSI/IEEE Standard 802.11*, 1999.
- [2] D. N. C. Tse and S. V. Hanly, "Linear Multiuser Receivers: Effective Interference, Effective Bandwidth and User Capacity," *IEEE Trans. Inf. Theory*, vol. 45, no. 2, pp. 641–657, Mar. 1999.
- [3] B. Lu, X. Wang, and J. Zhang, "Throughput of CDMA data networks with multiuser detection, ARQ, and packet combining," *IEEE Trans. Wireless Commun.*, vol. 52, no. 5, pp. 811–822, May 2004.
- [4] L. Tong, Q. Zhao, and G. Mergen, "Multipacket reception in random access wireless networks: from signal processing to optimal medium access control," *IEEE Commun. Mag.*, vol. 39, no. 11, pp. 108–112, Nov. 2001.
- [5] P. W. Wolniansky, G. J. Foschini, G. D. Golden, and R. A. Valenzuela, "V-BLAST: an architecture for realizing very high data rates over the rich-scattering wireless channel," in *Proc. of IEEE ISSSE*, Pisa, Italy, Sep. 1998, pp. 295–300.
- [6] S. Sfar, R. D. Murch, and K. B. Letaief, "Layered space-time multiuser detection over wireless uplink systems," *IEEE Trans. Wireless Commun.*, vol. 2, no. 4, pp. 653–668, Jul. 2003.
- [7] M. Levorato, P. Casari, and M. Zorzi, "On the performance of access strategies for MIMO ad hoc networks," in *Proc. IEEE GlobeCom*, Nov. 2006, pp. 1–5.
- [8] M. Zorzi, J. Zeidler, A. Anderson, B. Rao, J. Proakis, A. L. Swindlehurst, M. Jensen, and S. Krishnamurthy, "Cross-layer issues in MAC protocol design for MIMO ad hoc networks," *IEEE Wireless Commun. Mag.*, vol. 13, no. 4, pp. 62–76, Aug. 2006.
- [9] K. Sundaresan, R. Sivakumar, M. Ingram, and T.-Y. Chang, "Medium access control in ad hoc networks with MIMO links: optimization considerations and algorithms," *IEEE Trans. Mobile Comput.*, vol. 3, no. 4, pp. 350–365, Oct. 2004.
- [10] Y.-H. Lin, T. Javidi, R. L. Cruz, and L. B. Milstein, "Distributed link scheduling, power control and routing for multi-hop wireless MIMO networks," in *Proc. 40th Asilomar Conf.*, Pacific Grove, CA, Nov. 2006, pp. 122–126.
- [11] J. Liu, T. Y. Park, Y. T. Hou, Y. Shi, and H. D. Sherali, "Cross-layer optimization on routing and power control of MIMO ad hoc networks," Virginia Institute of Technology, Tech. Rep., Nov. 2006.
- [12] S. Cui and A. J. Goldsmith, "Energy efficient routing based on cooperative MIMO techniques," in *Proc. IEEE ICASSP*, Philadelphia, PA, Mar. 2005, pp. 805–808.
- [13] B. Hamdaoui and K. G. Shin, "Characterization and analysis of multi-hop wireless MIMO network throughput," in *Proc. ACM MobiHoc*, Montreal, Canada, Sep. 2007, pp. 120–129.
- [14] M. Hu and J. Zhang, "MIMO ad hoc networks: medium access control, saturation throughput, and optimal hop distance," *Journ. of Commun. and Networks*, pp. 317–330, Dec. 2004.
- [15] H. Jafarkhani, *Space-Time Coding: Theory and Practice*. Cambridge University Press, Sep. 2005.
- [16] M. Levorato, S. Tomasin, P. Casari, and M. Zorzi, "Physical layer approximations for cross-layer performance analysis in MIMO-BLAST ad hoc networks," *IEEE Trans. Wireless Commun.*, vol. 6, no. 11, pp. 4390–4400, Dec. 2007.
- [17] P. Casari, M. Levorato, and M. Zorzi, "Cross-layer design of MIMO ad hoc networks with layered multiuser detection," *IEEE Trans. Wireless Commun.*, submitted.
- [18] L. Muscariello, M. Mellia, M. Meo, M. A. Marsan, and R. L. Cigno, "Markov models of internet traffic and a new hierarchical mppp model," *Computer Communications*, vol. 28, no. 16, pp. 1835–1851, Sep. 2005.