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

☒ Reprint (Orig + 2 copies)
☐ Manuscript (1 copy)

☐ Technical Report (Orig + 2 copies)
☐ Final Progress Report (Orig + 2 copies)
☐ Related Materials, Abstracts, Theses (1 copy)

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

REPORT TITLE: Marco Levorato, Paolo Casari, Michele Zorzi, "On the Performance of Access Strategies for MIMO Ad Hoc Networks", IEEE Global Telecommunications Conference (GLOBECOM), San Francisco, CA, USA, 27 November - 1 December 2006.

is forwarded for your information.

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

Sincerely,

Michele Zorzi
University of California, San Diego

Enclosure 3

REPORT DOCUMENTATION PAGE	Form Approved OMB NO. 0704-0188

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE September 2006	3. REPORT TYPE AND DATES COVERED **Reprint** **01 June 2004 - 30 September 2006**
4. TITLE AND SUBTITLE On the Performance of Access Strategies for MIMO Ad Hoc Networks **Error! Not a valid bookmark self-reference.**		5. FUNDING NUMBERS **W911NF0410224**
6. AUTHOR(S) Marco Levorato, Paolo Casari, Michele Zorzi		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **University of California. San Diego** **9500 Gilman Drive, La Jolla, CA 92093-0407**		8. PERFORMING ORGANIZATION 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		10. SPONSORING / MONITORING AGENCY REPORT NUMBER **N/A**

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 address the impact of different access strategies in ad hoc networks with multiple antennas and MIMO communications. We employ a cross–layer designed MAC protocol that allows both for multiple simultaneous access to the radio medium and for proper exploitation of multiuser detection at the receiver for interference cancellation purposes.

Still, the network is subject to early deadlock if no access strategy is employed that reduces transmission persistency. To this aim, we study two types of exponential backoff, namely node–wise and destination–wise, and combine them with a form of cooperative agreement on who takes the role of transmitter or receiver.

The impact of all schemes is assessed and a comparison is pursued, focusing on typical network metrics (such as throughput and transmission delay among others) and showing when and why one technique performs better than the others.

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. 239-18
298-102

On the Performance of Access Strategies for MIMO Ad Hoc Networks

Marco Levorato*, **Paolo Casari***, **Michele Zorzi***†

*Department of Information Engineering — University of Padova — Via G. Gradenigo, 6/B, I–35131 Padova (PD), Italy
†Department of ECE — University of California at San Diego — 9500 Gilman Drive, La Jolla, CA 92093-0407
E–mail: `{firstname.lastname}@dei.unipd.it`

***Abstract*— In this paper, we address the impact of different access strategies in ad hoc networks with multiple antennas and MIMO communications. We employ a cross–layer designed MAC protocol that allows both for multiple simultaneous access to the radio medium and for proper exploitation of multiuser detection at the receiver for interference cancellation purposes.**

Still, the network is subject to early deadlock if no access strategy is employed that reduces transmission persistency. To this aim, we study two types of exponential backoff, namely node–wise and destination–wise, and combine them with a form of cooperative agreement on who takes the role of transmitter or receiver.

The impact of all schemes is assessed and a comparison is pursued, focusing on typical network metrics (such as throughput and transmission delay among others) and showing when and why one technique performs better than the others.

I. INTRODUCTION

AD HOC wireless networks have recently gained momentum as a means of introducing seamless access to network services, regardless of physical location, node mobility, and absence of a dedicated infrastructure. Well known standards have been deployed over the last decade [1], [2] in order to embody a comprehensive set of basic access rules and behaviors into a single set of protocols. To this extent, 802.11 [1] includes, for instance, a whole set of physical layer parameters that allow for different raw bit rates, and three access policies, relying on simple preemptive backoff, polling, or signaling message exchange in order to gain medium access.

Ad hoc networks are designed for flexibility and easy reconfigurability, trying to bridge the gap between user mobility and broadband wireless service access. Anyway, unremovable limits due to the very radio links used constrain the achievable performance, so that the available throughput for a single user typically tends to decay with the number of users [3], no matter how efficient the employed access scheme is.

Recent proposals for the application of directional or multiple antennas to wireless terminals have pointed out the benefits gained with these technologies in counteracting multiple access impairments. Consequences on MAC protocols have been described, e.g., in [4]–[6].

Multiple antennas may also be arranged to form a Multiple–Input–Multiple–Output (MIMO) system between transmitters and receivers, where ideally spatial (i.e., between pairs of antennas) uncorrelation is harnessed for encoding data in space and time, following a predefined structure (called a Space–Time Code, STC), or by multiplexing multiple data streams in space, transmitting them with different antennas simultaneously (Spatial Multiplexing, SM). For a comprehensive review of MIMO systems and STCs, the reader is referred to [7].

It is anyway not trivial to match MIMO systems at the physical layer with MAC protocols and link load control. To acquire the best advantage from the use of such techniques, one should design the whole system such that PHY and MAC features are jointly accounted for, while at the same time also trying to optimize relevant network metrics [8].

Ad hoc networks are inherently multi–access networks, where scarce radio resources are contended for by a given set of users, and simply "plugging" a new, even more powerful PHY layer to a previously existing MAC does not necessarily turn into significant networking improvements. The best design paradigm is therefore a cross–layer one, whereby the joint consideration of a greater set of degrees of freedom (both at PHY and MAC) can give much better performance than existing protocols.

In this work, we provide some insight on the effects of such a design perspective on ad hoc networks with MIMO communications, when random or distributedly controlled access strategies are taken into account. In particular, we discuss the pros and cons of random backoff techniques at the MAC layer, when compared to a different scheme which tries to achieve a distributed coordination among receivers and transmitters. Moreover, we study some interesting network metrics under the variation of appropriate access parameters.

II. SYSTEM AND NETWORK DESCRIPTION

A. Physical Layer Model and Approximation

In the following an outline of the physical layer model we assume for communications is provided. For a more complete description of our system, please refer to [9].

We assume that each node has multiple antennas available, which are exploited through SM with the double aim of $i)$ increasing raw bit rate per link, and $ii)$ combating interference through estimation and cancellation of unwanted incoming streams[1] (e.g., streams belonging to other links, that would collide if cancellation techniques were not applied). We suppose that the network is deployed inside a sufficiently rich scattering environment, where SM can be feasibly used by allowing each node to choose how many antennas to use for its transmission, following a certain rule.

The receivers apply a LAyered Space–Time MultiUser Detection technique (LAST–MUD) to separate upcoming signals, which is based on successive detection and cancellation of superimposed symbols, equivalent to the well known V–BLAST [10]. LAST–MUD enables receivers to detect more than one wanted data stream, while protecting them from interference, which is represented by superimposed unwanted data, removed from the useful contribution by generalized zero forcing and successive cancellation of detected signals.

Since channel estimation is a key phase of the detection process, each stream contains a training sequence that is used to acquire channel state information at the receiver. The rich scattering assumed allows to model channel effects as

---

[1] By stream we mean a fixed-size block of bits to be transmitted with one antenna, e.g., obtained by fragmenting a bigger data unit.

1-4244-0357-X/06/$20.00 ©2006 IEEE

Rayleigh. Due to complexity reduction or energy saving issues, or even to network effects, receivers may decide not to estimate channel conditions of all streams, leaving some as unestimated interference, which will then deteriorate the SINR of any other signal during the detection process.

As a final observation, we employ here an analytical approximation that models the effect of successive detection errors in the SINR as a Gaussian contribution. The details of this computationally efficient method are explained in [9], along with a more thorough description of LAST–MUD.

B. Network Protocol and Link Load Control Features

Different from traditional narrowband multiple access techniques, which try to avoid simultaneous transmissions, in our scheme the use of LAST-MUD makes it desirable to promote transmission parallelism. This makes scheduling significantly more challenging, since instead of just avoiding collisions the protocol must try to balance the channel usage in a distributed way, in order to maximize throughput while avoiding interference overload at the receivers.

As in 802.11 DCF [1], we use signaling RTS (CTS) messages to request (grant) transmissions, but insert there additional information about the requested (granted) number of antennas to be used on a given link. In fact, with SM, collisions between simultaneously sent packets may be solved by layered multiuser detection.

More specifically, RTSs are sent by nodes requiring link setup, specifying the requested number of antennas in the TX–RX MIMO link. In order to exploit SM to the maximum extent, multiple link requests may be inserted into the RTS, each with one or more antennas.

Receivers then decide how to grant transmissions to senders, controlling the link load, while trading off throughput for interference protection. The used *Follow Traffic* (FT) policy works under the constraint that at most N_s channels can be simultaneously estimated to separate incoming streams, each by means of a different training sequence. FT ensures adequate protection from interference by devoting at least $N_s/2$ training sequences to estimation and cancellation of interfering streams. The other $N_s/2$ degrees of freedom are used for the detection of the highest-power signals, with the only constraint that at least one intended packet should be included. Once receivers have made a decision about how to allocate the training sequences (i.e., which signals to detect) they respond with a CTS which contains a list of intended signals that are cleared for transmission.

Our algorithm works in synchronous frames. Each frame is composed of an RTS phase, where *all* transmitters dispatch their requests, a CTS phase where all receivers send back their grants, a data and an ACK phase. This may seem a major assumption, but in fact it is not, since frame synchronization is currently attainable in available standards [1]; moreover, results show that the cost of synchronization is justified by the very large improvements in terms of network performance.

This MAC behavior of the node still lacks an important design choice, i.e., how to reduce node persistency in transmission attempts, which could rapidly lead to deadlock if not properly controlled. A comparison between a random (backoff–based) and a distributedly coordinated control of transmissions is the main focus of this paper and will be addressed in the following Section.

III. ACCESS METHODS FOR MIMO AD HOC NETWORKS

As stated in Section II, the LAST–MUD algorithm at the receiver is subject to two main problems: the limited number of obtainable channel estimates, which is due to intrinsic limitations of the processing circuitry, and the need not to be overloaded by too many superimposed streams, which could experience too much mutual uncancelable interference. For these reasons, it is paramount that the receiver does not become overwhelmed by demultiplexing requirements.

As an immediate consequence, nodes cannot afford to access the radio channel in a completely random manner, whereupon traffic is injected in the network whenever there is need to transmit, since this would quickly lead the network into congestion. The devised RTS/CTS based handshake moves some steps in this direction, but does not solve the problem completely, as will be clear from the results in Section IV. Hence, we still need a means of forcing nodes to defer transmissions when temporary local congestion occurs; also, we should be able to guarantee that when excess traffic is offered, the applied policy does not lead to network deadlock, but is able to let some packets through.

As a first observation, note that carrier sense is not of much use here. A MIMO network should be designed to work with many nodes sending or receiving data at the same time, so that it would be disadvantageous to refrain from transmissions upon sensing a busy channel, because it is very likely (and in fact desirable) that the medium is busy most of the time. Access strategies should be designed keeping in mind i) the coexistence of multiple, potentially multi–antenna links, which follows from the previously described network structure; ii) the tradeoff between per–transmission SM (thus, raw transmission bit rate) and achievable throughput: as a rule of thumb, a weak SM gives higher per antenna transmission power and lower receiver load, but also low throughput; on the contrary, a more aggressive SM results in lower per antenna power and higher receiver load, increasing mutual interference between streams and lowering success ratio, and hence throughput. Finally, iii) a link load management technique is implemented in our scenario, that limits the number of grants inserted in the CTS whenever too many requests are sent by the transmitters. Note however that this policy is distributed, and therefore it cannot give any hard guarantee against congestion, but only steer system operation towards good working points.

A channel access able to maximize delivered traffic must take into account a means of reducing transmission persistency and balancing receivers and transmitters. Our main contribution here is to exploit three different policies for this purpose, assess their performance and explain why their behavior differs under certain network conditions. We consider two versions of exponential backoff and a form of distributed coordination, as explained in the following paragraphs.

A. Exponential Node–wise and Destination–wise Backoff

The node–wise version of exponential backoff in our context is administered as follows. Whenever a CTS is not received in response to an RTS, then the transmitter is forced to defer any other communication toward any other terminal for a certain time. The silence interval length is chosen to be a number of frames randomly distributed between 1 and an upper value BO_{max}, which is increased upon any subsequently failed attempt, following the law $BO_{max} = W \cdot 2^{N_u - 1}$, N_u being the number of unanswered RTSs. The parameter W represents the initial backoff window length, and can be set to change the intensity of retransmission attempts. We will refer to node–wise backoff as node–lock (NL). The maximum value reachable by BO_{max} is bounded by $W \cdot 2^5$.

The destination–wise version works similarly, but instead of blocking transmissions toward any node, it simply defers attempts toward the receiver which did not reply to a sent RTS. We expect this policy to allow for a higher throughput and lower packet transmission delay in the presence of equal window lengths, but to be more prone to congestion and LAST–MUD overload effects. This policy will be referred to in the following as dest–lock (DL).

B. Distributed Coordination using ACKs

Unlike exponential backoff, which is totally random, distributed coordination is based on the idea that, by achieving some agreement about which nodes should transmit or receive in a given frame, link setup could be easier, and it would be more likely to find the intended receivers free to listen. This accomplishment should anyway be reached at a minimum further cost in terms of increased signaling, so that efficiency is not compromised.

Following these guidelines, we recall that all receivers are supposed to send out ACKs to confirm correct reception, so that they could reserve nodes as destinations by including in the message an explicit request for them to become receivers in the following frame. This has the double advantage to piggy–back information into already existing signaling packets and to address nodes that have been idle or have transmitted, and are thus listening during the ACK phase of a frame. Reserving ACKs are priority messages. If one is received, the node refrains from transmission in the following frame, and waits for RTSs.

More in detail, each receiver wanting to reserve nodes in the following frame scans the first D packets of its queue and, with probability p_r, compiles a field of the ACK with the D corresponding destinations. Nodes recognizing themselves as reserved terminals refrain from transmitting in the following frame and wait for RTSs. With this scheme, idle nodes can be reserved as well. If not reserved, these idle nodes may decide to transmit in the following frame (with probability p_t), and preferably send packets to reserved nodes (which they know due to ACK overhearing). This access scheme will be referred to in the following as DSMA, an acronym for Distributed Scheduling for MIMO Ad hoc networks. A more thorough description and evaluation of DSMA can be found in [11].

We stress that this policy makes nodes alternate transmitter and receiver roles, because a node which was a receiver in a given frame cannot be reserved to be a receiver again in the following frame (sending its own ACK makes it deaf to any incoming reserving ACKs from other nodes). Moreover, this scheme partially solves some problems due to unestimated interference, that arise with distributed exchange of signaling messages. More details will be given in Section IV.

IV. SIMULATION RESULTS

A. Simulation Environment

To test the described policy, we have built a MATLAB simulator that accounts for the PHY communication level by means of the analytical approach devised in [9]. We have displaced a total of 25 nodes in a square grid inside a 100 m × 100 m area, each node having nearest neighbors at a distance of 25 m. Nodes are equipped with 8 antennas and transmit in the 5.8 GHz ISM band using BPSK. Transmissions take place in frames composed of RTSs, CTSs, DATA and ACKs. Each signalling message is 200 bits long. During the data phase, senders may use more than one antenna for communications, transmitting one stream of 1000 bits per used antenna. In this case, power is equally split among the active antennas, due to the lack of channel state information at the transmitter. Traffic is generated according to a Poisson process with rate λ packets per node per second. Each packet is randomly assigned a source and a length equal to $k \times 1000$ bits, with k uniformly chosen in $\{1, 2, 3, 4\}$. Unsent packets are kept by their sources in a backlog queue, which has a maximum length of 120 1000-bit streams. Under this configuration, the network is completely connected (all nodes are inside coverage range of each other). This choice is motivated by our goal to understand MAC issues. Extension to multihop scenarios is left for future study.

B. Results on Access Strategies

With our simulator, we have obtained the results shown in Figs. 1 to 5, where we have compared the behavior and performance of node–lock (NL), dest–lock (DL), and distributed scheduling (DSMA) using different metrics. For each strategy, the value of the initial window W or the pair (p_r, p_t) is given, as appropriate.

First, consider Fig. 1, depicting average throughput (defined as the average number of 1000-bit streams that are successfully received by their intended destinations per frame).[2] DL and NL have different behaviors for varying W. In particular, DL is a more aggressive policy. It allows nodes to send out more requests by just blocking single unavailable destinations. As a consequence, DL performs better than NL only if W is sufficiently high, such that congestion does not occur. For example, for $W = 1, 2, 4, 8$, DL is subject to a decay in throughput performance which is progressively mitigated by increasing W. This decay is mainly caused by the unsustainable amount of traffic generated due to node persistency in transmission attempts which eventually overloads the receiving stage and prevents a correct detection. Conversely, $W = 12, 16$ force longer silences on average, hence it is more likely that receivers become less loaded.

NL, on the other hand, imposes to defer any communication, having any transmitter turn into an available receiver for a given time upon any failure. Anyway, if W is too large the throughput saturates to a suboptimal value. With sufficiently low W, instead, NL outperforms the best throughput reached by DL.

Even if outperformed by NL from a throughput point of view, DL is very useful for keeping transmission delay (defined as the number of frames from packet generation to the packet transmission that ends correctly) as low as possible. Fig. 2 details this fact, which is a direct consequence of DL's aggressiveness. With DL, nodes can transmit more often, so that in low traffic scenarios they still experience a fair stream success probability with lower delay. Fig. 3, depicting the ratio of the correctly received 1000-bit streams to those sent, supports this deduction. Such considerations suggest that DL be used when low traffic is expected, while switching to NL at higher traffic and using, e.g., the average experienced delay as a measure of local network congestion for deciding when to switch from DL to NL. Studying and designing adaptive protocols is out of the scope of this paper, and is currently being addressed.

---

[2] Note that, with 802.11, the maximum throughput attainable in a completely connected network cannot exceed 1 stream per frame.

1-4244-0357-X/06/$20.00 ©2006 IEEE

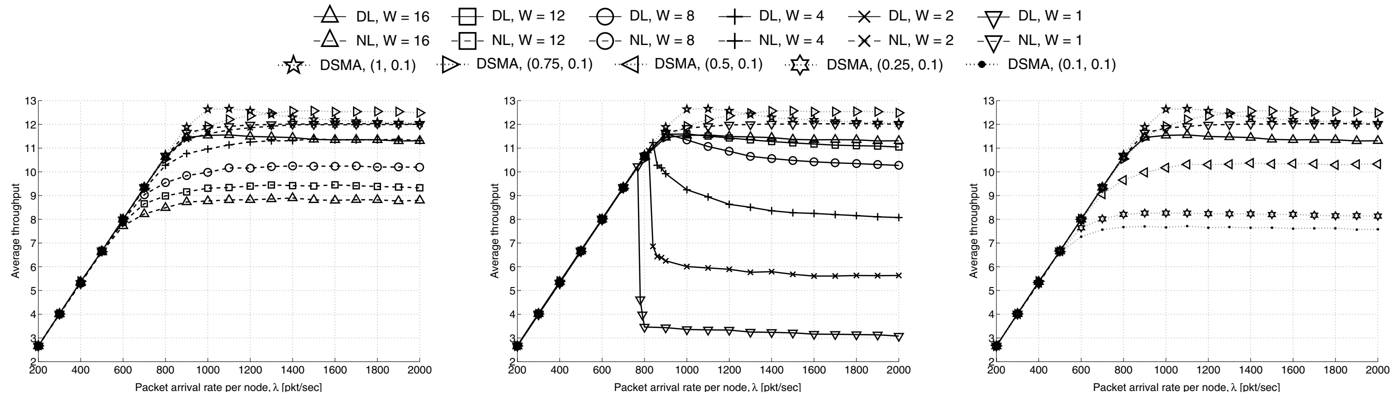

(a) Node–wise backoff (NL) (b) Destination–wise backoff (DL) (c) Distributed scheduling (DSMA)

Fig. 1. Average throughput as a function of traffic. Each graph reports one policy, along with the best results obtained with others.

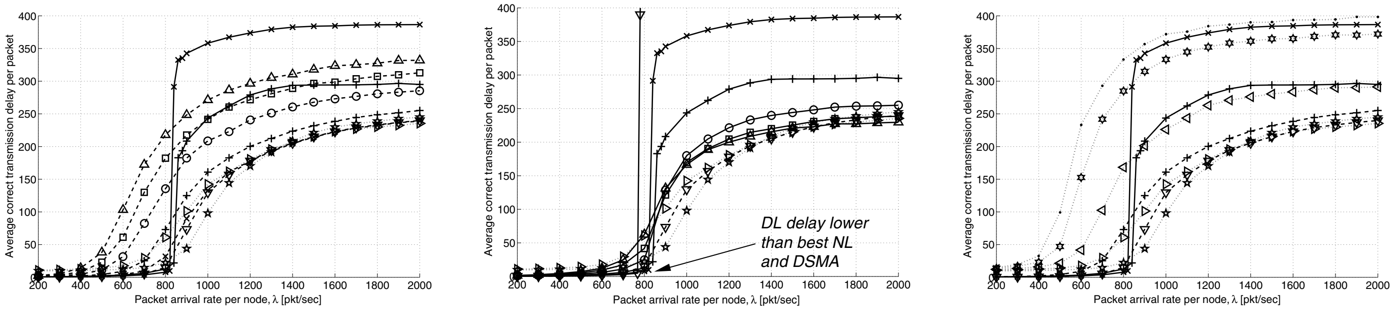

(a) Node–wise backoff (NL) (b) Destination–wise backoff (DL) (c) Distributed scheduling (DSMA)

Fig. 2. Average delay before correct transmission as a function of traffic. Each graph reports one policy, along with the best results obtained with others.

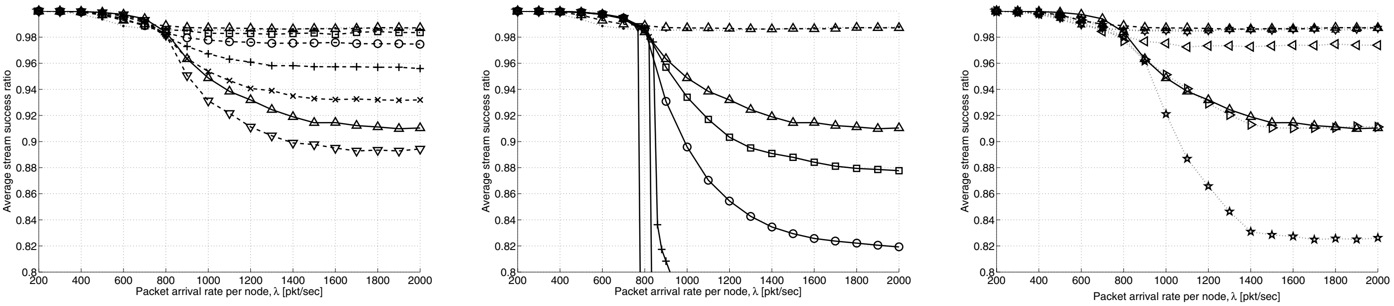

(a) Node–wise backoff (NL) (b) Destination–wise backoff (DL) (c) Distributed scheduling (DSMA)

Fig. 3. Average stream success ratio as a function of traffic. Each graph reports one policy, along with the best results obtained with others.

The behavior of distributed scheduling (DSMA) is different. By acquiring some coordination among receivers and transmitters, DSMA is able to reach a higher throughput, outperforming both backoff schemes. The two parameters p_t and p_r defined in Section III-B are crucial in determining the protocol behavior, and have been chosen so that an appreciable difference among curves is observed in the given results. As seen from Figs. 1 and 2, DSMA outperforms both NL and DL even in strong traffic scenarios, provided that p_r is sufficiently high, i.e., that an adequate number of reservation messages are sent by receivers. In particular, for $p_r = 0.75$ and $p_r = 1$, the network is pushed to a throughput of as much as 12.8 streams per frame. We also observe that for $p_r = 0.75$, the throughput increases more gradually but then achieves the highest value among all strategies. This interesting result shows that relying only on distributed scheduling of transmissions and receptions is not the most advantageous choice; the best performance is instead achieved by preserving a certain degree of randomness.

To obtain further insights on the behavior and applicability of the schemes presented, we depict in Figs. 4 and 5 the average number of links activated per node per frame, and the average number of transmitters per frame, respectively. By "link," we mean a node–to–node connection, regardless of the amount of spatial multiplexing used. In Fig. 5, only three curves per policy are reported. With these figures, it is possible to understand whether high throughput strategies prefer to load single connections with many streams, or to create multiple links each with smaller SM. In the former case, for example, the policy would prove to be more suited to delay–constrained connection–based networking, where it is important to convey high traffic on a given link (e.g., as part of a longer multihop path toward a final recipient). In the latter, the policy would be more applicable to information distribution scenarios, where a single source may want to address several destinations in order to spread traffic faster.

DL tends to allow more transmitters than NL (Fig. 5) and

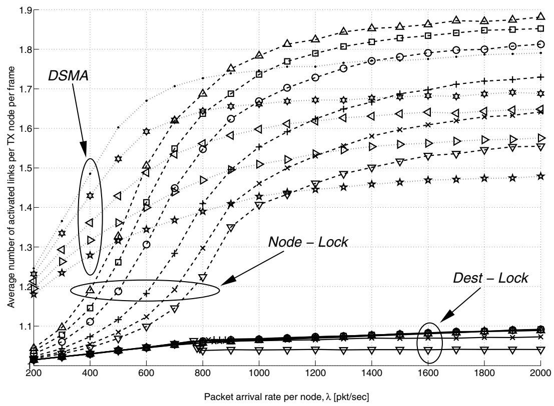

Fig. 4. Average no. of links per TX node per frame as a function of traffic.

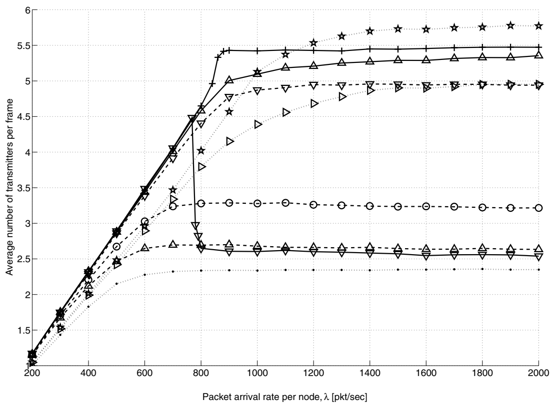

Fig. 5. Average no. of transmitters per frame as a function of traffic.

correspondingly more one–to–one connections (Fig. 4), with each connection having stronger spatial multiplexing, and thus higher throughput. Conversely, even the most permissive NL policy (for $W = 1$) enables a lower number of transmitters, each likely to connect to more than one receiver. NL thus achieves a lower data rate per link, but an overall better aggregate throughput.

DSMA instead tends to create multiple links originated from the same transmitter. As a result, DSMA curves in Fig. 4 all start from greater values than both backoff techniques, and then saturate to a maximum that is mainly determined by p_r (the probability that a receiver decides to make a reservation). Correspondingly, the higher the number of simultaneous links created, the lower the number of transmitters allowed, as seen from Fig. 5. This behavior results from DSMA design, whereby reservations are used to cycle the transmit role among nodes. If a receiver wants to transmit in the following frame, it may ask for up to D listeners each time. This engages one–to–many links with high probability, and is exploitable whenever a high degree of parallelism is needed along with higher bit rates. Such a scenario is found, e.g., in multihop wireless networks with many coexisting multimedia connections, where parallelism ensures a fast data spreading, and transmit role cycling helps mitigating starvation (the time before a receiving node can become a transmitter again).

As a final remark, consider again Figs. 1 and 3. We note that DL and NL experience high throughput in correspondence of a success ratio near 99%, whereas the max throughput DSMA configuration $(0.75, 0.1)$ undergoes 90% success only. This non-trivial result is explained by the different way DSMA organizes transmission parallelism. Namely, as many transmitters exist as in the best NL policy ($W = 1$), but each has slightly more active links per frame and with stronger SM than NL; moreover, due to the reservation mechanism, intended destinations are more easily reached than in random backoff policies. This reorganizes traffic in a sufficiently distributed fashion, such that success probability is still acceptable and performance is globally better. To sum up, DSMA allows for more SM traffic and enforces parallel one–to–many communications still achieving satisfactory success rates and the lowest average delay among all strategies. These are very important features for an ad hoc network where all nodes have heavy traffic requirements.

V. CONCLUSIONS

In this paper, we have compared three strategies for radio access in ad hoc networks with MIMO communications, namely Dest–Lock, Node–Lock and DSMA. We aimed first at showing the importance of correct transmission management in a network that relies on multiuser detection, since interference control is a primary issue to address.

Moreover, we have characterized all strategies under the point of view of throughput performance, success ratio, transmission parallelism, number of links per transmitter and delay of correct transmissions. We also highlighted relevant differences and analogies in how the network is driven by each strategy, and showed how spatial multiplexing and multiuser detection are used for enhancing throughput performance, and how transmitting nodes establish parallel links to different destinations. We have also studied how some relevant parameters affect the behavior of each policy, obtaining insights on how they could be tailored to match different traffic needs.

Future work on this topic includes a wider behavior characterization under different types of traffic, and the extension of our design to multihop scenarios.

ACKNOWLEDGMENT

This work is supported in part by the U. S. Army Research Office under MURI grant # W911NF-04-1-0224.

REFERENCES

[1] IEEE Standards Department, *IEEE Standard 802.11*. IEEE Press, 1999.

[2] ETSI BRAN, *HIPERLAN/2 Standard*. [Available online], http://portal.etsi.org/bran/kta/Hiperlan/hiperlan2.asp.

[3] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inform. Theory*, vol. 46, no. 2, pp. 388–404, Mar. 2000.

[4] A. Nasipuri, S. Ye, J. You, and R. E. Hiromoto, "A MAC protocol for mobile ad hoc networks using directional antennas," *Proc. of IEEE WCNC*, vol. 2, pp. 1214–1219, Sept. 2000.

[5] T. Korakis, G. Jakllari, and L. Tassiulas, "A MAC protocol for full exploitation of directional antennas in ad hoc wireless networks," *Proc. of the 4th ACM MOBIHOC*, 2003.

[6] R. Ramanathan, "On the performance of ad hoc networks with beamforming antennas," *Proc. of the 2nd ACM MOBIHOC*, 2001.

[7] H. Jafarkhani, *Space–Time Coding: Theory and Practice*. Cambridge University Press, Sept. 2005.

[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] M. Levorato, S. Tomasin, P. Casari, and M. Zorzi, "An approximate approach for layered space-time multiuser detection performance and its application to MIMO ad hoc networks," in *Proc. IEEE ICC*, Istanbul, Turkey, June 2006.

[10] G. J. Foschini, "Layered space-time architecture for wireless communication in a fading environment when using multiple antennas," in *Bell Labs Tech. J.*, vol. 1, 1996, pp. 41–59.

[11] P. Casari, M. Levorato, and M. Zorzi, "DSMA: an access method for MIMO ad hoc networks based on distributed scheduling," in *Proc. of ACM IWCMC*, Vancouver, Canada, July 2006.

1-4244-0357-X/06/$20.00 ©2006 IEEE