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) ☐ Technical Report (Orig + 2 copies)

☐ Manuscript (1 copy) ☐ Final Progress Report (Orig + 2 copies)

☐ Related Materials, Abstracts, Theses (1 copy)

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

REPORT TITLE: Marco Levorato, Stefano Tomasin, Paolo Casari, Michele Zorzi, "An Approximate Approach for Layered Space–Time Multiuser Detection Performance and its Application to MIMO Ad Hoc Networks", IEEE International Conference on Communications, (ICC), Istanbul, Turkey, Jun. 11-15, 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

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 May 2006	3. REPORT TYPE AND DATES COVERED **Reprint** **01 June 2004 - 31 May 2006**

4. TITLE AND SUBTITLE An Approximate Approach for Layered Space–Time Multiuser Detection Performance and its Application to MIMO Ad Hoc Networks	5. FUNDING NUMBERS **W911NF0410224**
6. AUTHOR(S) Marco Levorato, Stefano Tomasin, 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 consider a layered space–time multiuser detection technique and propose an analytical approximation for its performance. Our work is useful in two different stages of network design. On one hand, the approximation may be used to evaluate the bit and packet error performance of communications among a group of terminals making use of multiuser detection. On the other hand, it can also be seen as a very valuable tool from the networking point of view, as it may help in designing radio access control protocols based on multiuser detection, as analytical formulas are very fast to evaluate, in contrast to bit–level simulations that may need a long time to complete. Thus, an analytical formulation is important, because it helps discriminating among different protocol alternatives, speeding up considerably the protocol design phase and the development of new radio access policies for multiuser networks.

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**

An Approximate Approach for Layered Space–Time Multiuser Detection Performance and its Application to MIMO Ad Hoc Networks

Marco Levorato⋆, Stefano Tomasin⋆, Paolo Casari⋆, Michele Zorzi⋆†
⋆DEI — 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 consider a layered space–time multiuser detection technique and propose an analytical approximation for its performance. Our work is useful in two different stages of network design. On one hand, the approximation may be used to evaluate the bit and packet error performance of communications among a group of terminals making use of multiuser detection. On the other hand, it can also be seen as a very valuable tool from the networking point of view, as it may help in designing radio access control protocols based on multiuser detection, as analytical formulas are very fast to evaluate, in contrast to bit–level simulations that may need a long time to complete. Thus, an analytical formulation is important, because it helps discriminating among different protocol alternatives, speeding up considerably the protocol design phase and the development of new radio access policies for multiuser networks.**

I. Introduction

SINCE the early pioneering studies on Multiple–Input–Multiple–Output (MIMO) techniques for radio communications [1], the research efforts in this field have grown at a very fast rate. The MIMO approach to the wireless communications scenario paves the way for a number of improvements, such as a great increase of the raw available bit rate in a single–link connection, interference suppression capabilities, new techniques for multiuser detection, etc. As per the first of these, which had proved to be a very challenging issue in radio digital transmission before the introduction of such technologies, MIMO has been defined as the "key" to enable gigabit–per–second wireless links [2].

MIMO systems are realized by embedding multiple antennas in communicating terminals. With more than one antenna available, each terminal may exploit an increased number of degrees of freedom, that in turn enable the previously listed advanced communication features. For instance, when establishing a link with another terminal, a transmitter could decide to split its data, forming a set of shorter streams. Each stream is associated to a different antenna, and then all streams are transmitted simultaneously — a technique known as *spatial multiplexing* [2]. Since each antenna transmits shorter sequences, the overall transmission takes a shorter time to complete, thereby allowing for higher bit rates. At the receiver side, this generates a spatial superposition of as many symbols per symbol time as the used transmit antennas, so that the receiver needs some algorithms that allow for symbol separation and detection. The well–known V–BLAST scheme [1], [3] is the first one designed for such a purpose. Other schemes rely on more advanced mathematical processing of the receiver antenna outputs [4] to detect symbols, and then cancel their interference over the other superimposed ones, so as to globally improve the probability of correct detection. These techniques require the presence of multiple antennas at the receiver: as a general rule, the higher the number of available antennas, the better the decoding performance.

Recently, there has been increasing interest in the application of MIMO techniques to *ad hoc networks*. Ad hoc networks are collections of terminals linked together by means of packet radio communications only, without any underlying infrastructure and with no need for supervision by an external network operator. This kind of network offers great flexibility during the deployment phase, allowing for a fast and cheap network setup, and may represent the only affordable solution in certain scenarios (such as battlefields, natural disaster recovery, data sharing during business meetings, etc.). Ad hoc networks feature a great ease of access (typically plug–and–play–style) and are a very likely means of offering advanced and more complex services to mobile users in the near future.

Ad hoc networks are impaired by some technical limitations that arise due to the wireless propagation of transmitted signals. A well–documented problem is for example the "hidden terminal" [5], [6], a primary cause of "collisions" at receiving stations: being unable to hear each other's transmission, two nodes may transmit to the same receiver simultaneously resulting in discarding both packets. Furthermore, radio links need to be operated at the proper data rate in order to be reliable. Even if the IEEE 802.11a protocol for wireless networks enables high data–rate communications, up to 54 Mbps [7], such speeds are achievable only at very short distances.

In order to alleviate the problems that arise in data communications over unreliable radio links, ad hoc nodes may be equipped with multiple antennas. The MIMO technology, in particular, may be applied to obtain more reliable links as well as a higher transmission rate at longer distances [2].

However, from the point of view of the mathematical characterization and simulation (and hence from the point of view of protocol design), the algorithms typically used as part of MIMO techniques are often mathematically quite involved, and it is all but simple to embed them in network analysis. Simulation is possible in principle, but when combining the bit–level time scale (needed to simulate MIMO receiver processing) with that on which network protocols typically operate, obtaining good statistical results becomes

1-4244-0355-3/06/$20.00 (c) 2006 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE ICC 2006 proceedings.

computationally very demanding.

In order to alleviate the task of reproducing every single MIMO transmission and reception step, two approaches may be followed. The first is to use approximations of the overall "gain" (over traditional links) that could be obtained when operating MIMO links. This is a viable solution, but may lead to some unwanted oversimplification and to results that may not be completely meaningful. The importance of a good and realistic physical model when assessing ad hoc networking issues has been stressed in the recent literature (e.g., see [8]). The second option is to develop approximate formulas for the whole process involving spatial superposition, interference cancellation, and any other MIMO feature that takes part in the symbol transmission and detection phases, rather than just reducing it to a simple global gain factor. Such formulas would have to be fast to evaluate, while being able to reproduce the real behavior of the implemented technology with sufficient precision.

To this extent, we build a novel analytical framework for the analysis of the layered space–time multiuser detection (LAST–MUD) technique proposed in [4]. Besides being a technical contribution in itself, this approximate approach is then used to derive network protocol performance results in a much faster way than would be allowed by bit–level simulation.

Our paper is organized as follows. A description of the system model is carried out in Section II and is followed by the developed analytical models in Section III. In Section IV, we compare the analytical results with the corresponding simulations of the described system. Finally, Section V concludes our paper.

II. System Description

Here we briefly summarize the decoding process at the physical layer and its mathematical model. Full details can be found in [4].

Consider a packet radio network in a rich scattering environment. Nodes in the network have multiple antennas, that are used to exploit the fading process and to improve the available bit rate by means of spatial multiplexing [2].

We assume that each user may choose a subset of antennas for spatially–superimposed transmissions (depending for instance on traffic needs), but constrain the number of transmitted bits *per antenna* in a packet to be a constant, say 1000 bits per antenna (each 1000-bit block transmitted by one antenna will be called "stream" in the following). This simplifying assumption is made here for ease of notation, but it is by no means critical and may be removed to take into account the simultaneous transmission of packets having different lengths.

Spatial multiplexing needs symbol–by–symbol decoding. In particular, the receiver separates the incoming signals by means of a decorrelating layered space–time signal processing technique [4]. Suppose that K transmitters, $k = 1, \dots, K$, are sending data to a receiving node, each transmitter employing u_k antennas. The total number of streams to be detected is given by $U = \sum_{k=1}^{K} u_k$. Then, let $\mathbf{b} = [b_1, \dots, b_U]^T$ be the U-length column vector where each element is a symbol coming from one of the U transmitting antennas, and ^{T} denotes transposition.

As signals propagate from the transmitting antennas to the A antennas of the receiver, they undergo fading effects and path losses. We assume frequency–flat Rayleigh fading as a good model for the rich scattering network environment cited before, so that channel gains may be described by an $A \times U$ matrix $\mathbf{H}$ with zero–mean circularly Gaussian complex entries $h_{a\ell}$. The variance of each entry is determined by the per–stream path loss.

Assuming perfect channel estimation and symbol synchronization, the receiver processes antenna outputs on a symbol time basis, in order to extract sufficient statistics that enable spatial de–multiplexing. Let $\mathbf{h}_a = [h_{a1}, \dots , h_{aU}]$ be the a-th row of matrix $\mathbf{H}$, and write the received signal as $r_a = \mathbf{h}_a \mathbf{b} + n_a$, where n_a is the noise sample at the a-th antenna, modeled by a zero–mean complex Gaussian random variable with variance σ^2. Note that the product $\mathbf{h}_a \mathbf{b}$ outputs a scalar. The received signal at antenna a can be written as

$$r_a = \mathbf{h}_a \mathbf{b} + \mathbf{h}_a^{int} \mathbf{b}^{int} + n_a \, , \tag{1}$$

where we also explicitly accounted for an unestimated interference term coming from symbols $\mathbf{b}^{int}$, with $\mathbf{h}_a^{int}$ representing associated channel gains as seen from antenna a. The sufficient statistics for extracting signals is then given by [4]

$$\mathbf{m} = \sum_{a=1}^{A} \mathbf{h}_a^{\dagger} r_a = \mathbf{R}\mathbf{b} + \mathbf{n} + \mathbf{i} \tag{2}$$

where $\mathbf{R} = \sum_{a=1}^{A} \mathbf{h}_a^{\dagger} \mathbf{h}_a$ is the $U \times U$ space cross–correlation matrix, $\mathbf{n} = \sum_{a=1}^{A} \mathbf{h}_a^{\dagger} n_a$ is the filtered Gaussian noise vector with covariance $\sigma^2 \mathbf{R}$, and † denotes complex transposition. Space–matched filtering also acts on the previously described unestimated interference, so that a contribution $\mathbf{i} = \sum_{a=1}^{A} \mathbf{h}_a^{\dagger} \mathbf{h}_a^{int} \mathbf{b}_{int}$ is present in $\mathbf{m}$. The explicit modeling of a pure interference term is important point in an ad hoc scenario, as terminals may have different capabilities, and therefore may only be able to decode a limited number of signals, or they may decide to reduce the processing complexity when possible (e.g., in order to save energy), neglecting low–power interference in the decoding process.

Furthermore, since processing real correlation matrices yields a decoding benefit in terms of higher effectiveness, we extract the real part of $\mathbf{R}$ prior to calculating the sufficient statistics. This forces to use real signal constellations. In the following, we shall therefore write $\mathbf{R}$, with the understanding that it means $\mathrm{Real}[\mathbf{R}]$.

The receiver successively performs the detection of the streams and, upon each detection, the interference due to the already detected users is removed from the received signal. The order in which users are detected and cancelled is based on their Signal–to–Noise Ratio (SNR). Let us indicate with $\mathbf{R}(i)$, $i = 1, \dots , U-1$ the $(U-i+1) \times (U-i+1)$ cross correlation matrix before the detection of the i-th ordered user. We also set $\mathbf{R}(1) = \mathbf{R}$ (as previously defined) for the first detection.

The user with the highest SNR at the i-th detection stage is selected for processing by the receiver. Assuming that the receiver does not know the statistics of $\mathbf{i}$, the user index is obtained as

$$k_i = \arg\max_j \frac{\sigma_{b_j}^2}{\sigma^2 \left[\mathbf{R}(i)^{+}\right]_{(j,j)}} \, , \tag{3}$$

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE ICC 2006 proceedings.

where $\mathbf{R}(i)^+$ is the Moore–Penrose pseudoinverse of $\mathbf{R}(i)$ [9], and $\sigma^2_{b_j} = \mathrm{E}[b_j^2]P_j$, with P_j representing the transmission power level of the j-th stream.

The detection of user k_i is performed by first weighing the elements of the sufficient statistics at step i, $\mathbf{m}(i)$, with a set of coefficients given by the k_i-th column of $\mathbf{R}(i)^+$, namely $\mathbf{w}(i) = [\mathbf{R}(i)^+]^{k_i}$. This choice of the weight factors minimizes the interference due to users $k_{i+1}, \dots, k_U$.

The sample $\tilde{b}_{k_i}$ is finally obtained as

$$\tilde{b}_{k_i} = \mathbf{w}(i)^\dagger \mathbf{m}(i) \tag{4}$$

and passed through a hard detector, which provides the symbol estimate $\hat{b}_{k_i}$.

Prior to the following detection step $(i+1)$, user k_i's interference is removed from the received signal and the sufficient statistics vector, yielding [4]

$$\begin{aligned} \mathbf{r}(i+1) &= \mathbf{r}(i) - \mathbf{H}^{k_i}\sqrt{P_{k_i}}\,\hat{b}_{k_i} \\ \mathbf{m}(i+1) &= \sum_{a=1}^{A} \mathbf{h}_a(i+1)^\dagger r_a(i+1) \end{aligned} \tag{5}$$

where $\mathbf{r} = [r_1, r_2, \dots, r_A]^T$, $\hat{b}_{k_i}$ is the hard–detected bit belonging to user k_i and $\mathbf{H}^{k_i}$ is the k_i-th column of the matrix $\mathbf{H}$. We set $\mathbf{r}(1) = \mathbf{r}$, $\mathbf{h}_a(1) = \mathbf{h}_a$, $\mathbf{m}(1) = \mathbf{m}$ and obtain the next step channel vectors $\mathbf{h}_a(i+1)$ by striking out the k_i-th entry of $\mathbf{h}_a(i)$ and the correlation matrix $\mathbf{R}(i+1)$ by striking out both the k_i-th row and the k_i-th column of $\mathbf{R}(i)$.

III. Performance analysis

A. Gaussian approximate analytical approach

In this analysis we include both the effects of imperfect detection and the interference coming from undetected nodes. Let $\mathcal{N}$ be the set of users to detect and denote the set of already detected users before stage i as $\mathcal{K}(i) = \{k_1, \dots, k_{i-1}\}$ and the set of unestimated interfering users as $\mathcal{I}$.

We model the residual interference due to previous errors in the detection process as an additive gaussian signal term, with zero mean and variance depending on the error rate. A similar assumption has been, e.g., in [10]. In particular, suppose that user k_i is affected by a bit error rate $P_e(k_i)$, which will be determined in the following. The error signal will be

$$e(k_i) = \begin{cases} 0, & \text{with probability } 1 - P_e(k_i) \\ -2b_{k_i}, & \text{with probability } P_e(k_i)/2 \\ +2b_{k_i}, & \text{with probability } P_e(k_i)/2 \end{cases} . \tag{6}$$

Hence, the variance of the interference signal which corresponds to user k_i can be written as

$$\sigma_e^2(k_i) = 4\sigma^2_{b_{k_i}} P_e(k_i) \,, \tag{7}$$

Let $\gamma(k_i)$ be the Signal to Noise and Interference Ratio (SNIR) in the detection of user k_i, $\mathbf{R}_{int} = \sum_{a=1}^{A} \mathbf{h}_a^\dagger \mathbf{h}_a^{int}$, and $[\,\cdot\,]^k$ be the k-th column of a matrix. The SNIR can be obtained as expressed in (10). There, the numerator accounts for the post–detection useful signal power, whereas in the denominator, the first term expresses the effects of space–matched filtering on thermal noise, the second is the power of unestimated interference terms, while the third and the fourth respectively account for average error signal power due to previous possibly wrong cancellations and for residual interference due to not yet canceled streams, which could arise if the receiver tries to separate more incoming streams than available receive antennas.

From the SNIR, the bit error rate can be obtained as

$$P_e(k_i) = Q\left(\sqrt{\gamma(k_i)}\right) , \tag{8}$$

where $Q(\cdot)$ is the complementary Gaussian distribution. For uncoded transmissions, the packet error rate can be easily derived with the expression

$$\mathrm{PER}(k_i) = 1 - [1 - P_e(k_i)]^{P_S} , \tag{9}$$

where P_S is the packet size in bits. Other expressions, though more involved, can be similarly used to compute the PER in the presence of coding.

B. Exhaustive analytical approach

We now develop an alternative way of computing the error probability that is meant to be exhaustive, yet more complex from a computational point of view.

We start from the correct detection probability of the first user. The first user is not affected by error propagation due to erroneous past detections, and its error probability depends only on the channel conditions and noise level. Let P_{e_1} be the error probability for the first user. For the second user, the error probability P_{e_2} is related both to the channel conditions and to errors in the detection of the first user. In order to estimate P_{e_2}, all possible configurations of wrong and correct detection are considered, and for each configuration, the SNIR is evaluated. The process of considering all possible error configurations and consequent error probabilities is repeated for all users. We obtain a tree where each node corresponds to an error configuration of detected users and there are three branches departing from each node, corresponding to the two error configurations of a BPSK symbol and the case of a correct decision.

In [11], we have developed a different approximation based on the preceding considerations. The aim was to obtain the system BER by resorting to the error tree exploration, but also to reduce the exponential complexity of this approach to a lower level, by pruning those subtrees where the number

$$\gamma(k_i) = \frac{|\mathbf{w}(i)^\dagger \mathbf{R}^{k_i}|^4 \sigma^2_{b_{k_i}}}{\frac{N_o}{2}||\mathbf{w}(i)^\dagger \mathbf{H}^\dagger||^2 + \sum_{k\in\mathcal{I}} |\mathbf{w}(i)^\dagger \mathbf{R}^k_{int}|^2 \sigma^2_{b_k} + \sum_{k\in\mathcal{K}(i)} |\mathbf{w}(i)^\dagger \mathbf{R}^k|^2 \sigma_e^2(k) + \sum_{k\in\mathcal{N}\setminus\mathcal{K}(i)} |\mathbf{w}(i)^\dagger \mathbf{R}^k|^2 \sigma^2_{b_k}} \tag{10}$$

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE ICC 2006 proceedings.

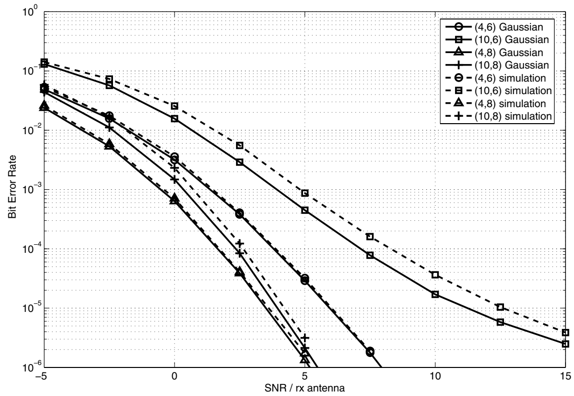

Fig. 1. Performance of the Gaussian approach (Section III-A) to LAST–MUD BER evaluation.

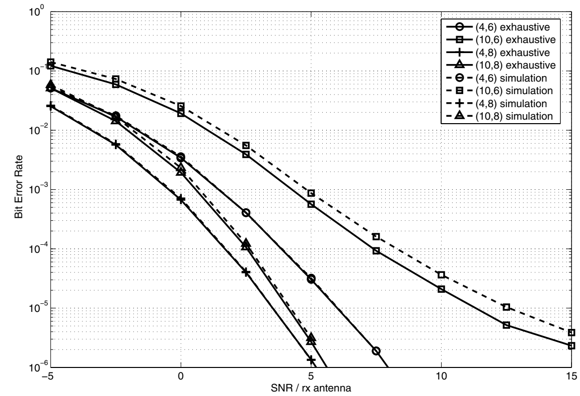

Fig. 2. Performance of the exhaustive approach (Section III-B) to LAST–MUD BER evaluation.

of errors exceeds a threshold (say, ℓ). In fact, if too many symbols are in error, we can assume that all subsequent detections have a very low SNR and hence an error probability close to 0.5. In [11], we set $\ell = 1$ and showed that, for network evaluation purposes, this solution gives very good BER assessment accuracy up to a certain extent, but tends to overestimate the PER in some network environments, when excess unknown interference is present. Later on (Section IV), we will instead show that the Gaussian approximation developed in Section III-A is more accurate for the evaluation of network performance metrics.

In the following, we will describe the BER prediction accuracy of the Gaussian approach and compare it with exhaustive error tree exploration for the sake of completeness.

C. Performance comparison of the approximate techniques

In Figs. 1 and 2 the Bit Error Rate approximation performance of each of the previously described techniques is depicted. There, the notation (T, R) means that R receive antennas are used to decode T spatially superimposed streams, each coming from a defferent user, with each user experiencing the same received SNR. Note that T may be greater than R: this fact is to be accounted for in ad hoc networks, where it is not known *a priori* how many transmissions would take place at a given instant. If it should happen, it is known that more than R incoming streams would make matrix $\mathbf{R}$ singular. The algorithm would then rely on the pseudo–inversion of the correlation matrix $\mathbf{R}(i)$ at each step, in order to obtain processing weights that maximize the detection reliability.

The approximate procedure reported in Section III-A proves to be very good for a small number of simultaneous incoming streams (i.e., 4), while it is less accurate if the number of streams to decode is increased to 10 (see Fig. 1). This difference is due to the model considered for the propagation of detection errors, which are approximated as a Gaussian contribution in order to be included in the SNIR expression (10). While this model is very accurate for a small number of streams, its accuracy degrades (though still being reasonably good) in the presence of a larger number of signals.

Fig. 2 shows the BER performance obtained via the exhaustive method described in Section III-B. In this case, the approximation is more precise as it is able to predict the BER accurately for all configurations except $(10, 6)$, that is the most difficult one from the point of view of error propagation modeling. Indeed, this approximation is closer than the other to the simulated curve.

Notice that the exhaustive method has a higher computational complexity, that grows exponentially with the number of streams, whereas the complexity of the simpler approximate approach is only linear. Also, the accuracy of the simpler scheme is already fairly good, and the small improvement given by the exhaustive search does not seem to justify the significant additional complexity. For this reason, we will use the simpler Gaussian approach in our network simulation, as described in the following Section.

IV. Network protocol performance results

In this Section, we discuss the application of the analytical approximation of Section III-A to the simulation of ad hoc networks with multiple antennas.

Our protocol is based on a preliminary exchange of control messages, which we call Request–to–Send (RTS) and Clear–to–Send (CTS), according to the common 802.11 terminology [7]. This exchange precedes the actual data packet transmission phase, which is finally concluded by an acknowledgement message (ACK). We organize the handshakes in frames composed of an RTS, a CTS, a data and an ACK phase. We require transmitting nodes to be synchronized, so that all packet transmissions start at the same instant within each phase. Even though frame–level synchronization may seem a major drawback of our approach, we stress that i) synchronization (e.g., slot by slot) is achievable and currently used, e.g., in protocols belonging to the 802.11 standards [7]; ii) in this example, we assume a completely connected network, where all nodes are within carrier sensing range of each other, so that synchronization is easier to achieve; finally, iii) this assumption could be relaxed if nodes hearing the start of unsynchronized transmissions leave these as part of the unknown interference term introduced in Eq. (2). Studying the impact of unsynchronized transmissions on the performance of this protocol is beyond the scope of this paper, and is currently being investigated.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE ICC 2006 proceedings.

A. Simulation Environment and MAC Policies

We assume a fully connected network of 25 nodes arranged on a 5×5 square grid and fix the distance between nearest neighbors to 25 m. Channel variations are assumed to be negligible for the duration of a frame, which is reasonable, e.g., when considering transmissions in the 5.8 GHz ISM frequency band and the use of BPSK with 99% in–band power, where the raw bit rate per used antenna is 7.5 Mbps. Also consider that when operating at 5.8 GHz, the inter–antenna spacing may be set to less than 3 cm, still preserving uncorrelation between copies of the same signals received by different antennas [12]. Thus, it would be possible to fit a linear array of 8 antennas in 21 cm (less than the width of a laptop screen). Hence, we assume that nodes have 8 antennas each.

Packets in the network are generated according to a Poisson process of parameter λ (packets per second per node): each packet is assigned a random destination and a random length of $k \times 1000$ bits, with k uniformly chosen in the set $\{1, 2, 3, 4\}$. Nodes can keep a maximum of 30 packets in their queue, regardless of the number of bits. Packet transmission is synchronized across all transmitters, and the packet duration is fixed to 200 bit intervals for signaling packets and to 1000 bit intervals for data packets. Note that spatial multiplexing can be used for data packets, so that the number of sent data bits is 1000 times the number of antennas used. All transmissions are uncoded. Preceding all transmitted packets is a training sequence for channel estimation, one per used tx antenna. Nodes are assumed to be able to apply the receive algorithm to no more than N_s signaling packets or N_s spatially multiplexed data streams (one stream per transmit antenna). In our results, we assume $N_s = 32$.[1]

Any transmitter collects packets from its queue on a FIFO basis, and composes RTS messages accordingly, including both the destination and the number of antennas to be used, so that the destination knows how many training sequences are needed to separate and decode all incoming streams. When sending back a CTS, the receiver is in charge of making the ultimate decision on the number of antennas to be used by the transmitter, so it includes this information into the message. How the decision is made depends on the prescribed policy, which is tunable as a design choice. We consider here a "Follow Traffic" (**FT**) policy, that is described below.

In order to preserve at least some amount of throughput, but to offer a sufficient level of protection against unwanted interference, each node saves $N_s/2$ sequences for cancellation purposes only, and manages the others as follows. First, it selects the wanted stream with maximum received power and grants it in the CTS, and then evaluates other requests in order of decreasing received power. If the node is the intended receiver as expressed in the request being evaluated, then it allows the transmission of the corresponding 1000-bit stream with the CTS, otherwise it simply reserves a training sequence for channel estimation and subsequent interference cancellation, with the aim of enhancing the correct decoding probability of the desired incoming transmissions. The final outcome of this policy is a CTS packet containing all nodes allowed to transmit and the number of antennas they are allowed to use (thus, the number of 1000-bit streams that are granted transmission).

As a CTS is received, nodes transmit data following the instructions contained inside it and wait for a stream–wise ACK from the receiver. A packet is not removed from the queue until all of its streams have been acknowledged. If a CTS should not be received for any reason, a backoff timer is set up, whose duration lasts a number of frames uniformly chosen between 1 and B_{max}. As the local network load is already kept under control by receivers using CTSs, backoff here is meant to be a means of de–synchronizing transmissions, so that it becomes less likely that too many nodes try to send data at the same time. More in detail, we set up two different backoff strategies. The first one blocks transmissions towards an intended receiver which did not reply with a CTS to a transmission request (Dest–Lock); the second strategy blocks transmissions toward any node upon any lack of CTS response (Node–Lock), and is similar to common backoff techniques used in MAC protocols for random access in wireless networks, such as the well known 802.11 [7].

If more than one attempt should fail, B_{max} is increased exponentially as a function of the number of failed attempts, using the rule $B_{max} = W \cdot 2^{N_{fail}-1}$, where W is a startup backoff window length which we set to 16. Eventually, the packet is dropped if too many failures occur.

For the simulation of such a network environment, we have developed a MATLAB simulator, where we implemented both the detailed bit–level processing of LAST–MUD and the approximate analytical procedure of Section III-A to evaluate the BER for spatially multiplexed transmission.

B. Networking results

With our simulator, we have obtained some results that are shown in Figs. 3 and 4. They respectively depict the average throughput (i.e., the number of successfully decoded streams per frame), and the average queue length measured in number of backlogged 1000-bit streams, respectively.

The accuracy of the analytical approach (recall that we used the approximation described in Section III-A) is confirmed in all figures, as simulation fits the analysis quite well, especially for reasonable values of the packet generation rate λ (the congestion region of the curves is of lesser interest).

For the sake of completeness, the figures also report the performance evaluation obtained through the pruned error tree approach mentioned in Section III-B, as a lower complexity variation of the exhaustive error tree exploration. The pruned tree approach is described more in–depth in [11]. In both cases, we observe that a better approximation accuracy is achieved by the Gaussian approach, which predicts the simulated network behavior very well in both the Node–Lock and the Dest–Lock case. The pruned tree approach, on the other hand, performs well at traffic values for which the network is still uncongested, but after that shows a different behavior than simulation predicts, especially in the Dest–Lock case. The reasons behind this kind of discrepancy are to be searched in the higher aggressiveness of destination–wise backoff, which tends to

[1] Although matrix $\mathbf{R}$ becomes singular for more than 8 incoming streams (with 8 receive antennas), a higher N_s is needed in order to counter some network effects that would lead to excess unknown interference. For a thorough description of these effects, see [11, Section III-B].

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE ICC 2006 proceedings.

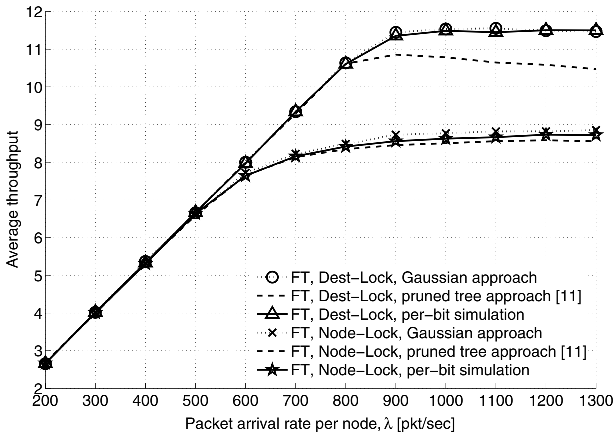

Fig. 3. Average throughput as a function of network load for different backoff strategies, simulation and analytical approaches.

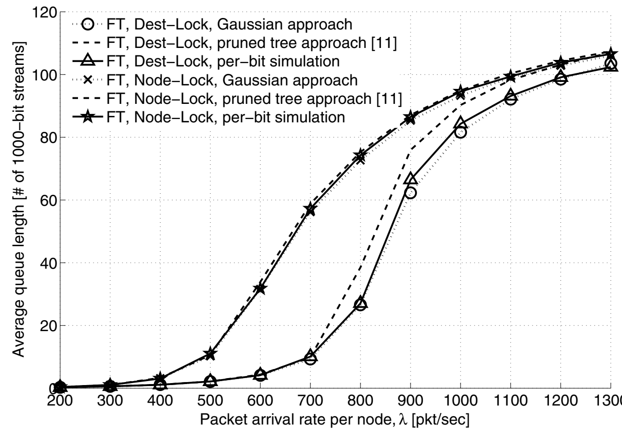

Fig. 4. Average queue length as a function of network load for different backoff strategies, simulation and analytical approaches.

leave nodes more free to transmit, thus increasing network utilization and average interference affecting transmissions in each frame. A thorough investigation of this behavior can be found in [11]. In this more crucial configuration, higher interference worsens the BER performance of the pruned tree approach, due to the thresholding that forces any symbol detection after the occurrence of errors to undergo a success probability of 0.5. The Gaussian approach circumvents this problem by “spreading” the propagation of decoding errors and using the modified SNIR in (10), which is a much softer metric than a fixed error chance, and is capable of predicting overall effects with a higher precision, on average.

Even if a comparison between backoff strategies is not the main objective here, we stress that the Node–Lock policy is more conservative than Dest–Lock, because it forces nodes to defer transmissions towards any other node upon any failure, while Dest–Lock focuses on a single receiver. This behavior explains why throughput reaches greater values for Dest–Lock (and queues are correspondingly left emptier). The purpose of these sample results is not to give a comprehensive study of the networking performance of our scheme, but rather to show that the analytical physical layer approximation proposed in this paper leads to accurate networking evaluations as well. This makes it possible to replace lengthy simulations on a bit–level time scale with simulations where the physical layer details are adequately captured by closed–form formulas, which lead to statistically meaningful networking results in a much shorter time.

V. Conclusions

In this paper we have presented and discussed techniques to model the behavior of layered space–time multiuser detection in ad hoc networks. We have evaluated their accuracy, and shown that the Gaussian approximation is nearly as precise as the exhaustive approach, and that both predict a bit error probability that is very close to the values obtained with a fully detailed bit level simulation.

Furthermore, we have applied the approximate technique to network performance evaluations and shown that the obtained results match simulations very well, but are much faster to obtain compared to complex and long bit–level MIMO systems simulations. Thus, we believe this to be a very helpful tool in network design, that provides a quick and accurate way to compare different protocol solutions.

Future work includes some additional refinements to our analytical approach and its extension to other MIMO schemes, as well as an extensive study of the MAC performance and of different options for protocol optimization in a MIMO ad hoc networking context.

Acknowledgment

This work is supported in part by the U. S. Army Research Office under the Multi–University Research Initiative (MURI) grant # W911NF–04–1–0224.

References

[1] 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.

[2] A. J. Paulraj, D. A. Gore, R. U. Nabar, and H. Bölcskei, “An Overview of MIMO Communications: A Key to Gigabit Wireless,” *Proc. IEEE*, vol. 92, no. 2, pp. 198–218, Feb. 2004.

[3] 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,” *Proc. of the ISSSE–98*, Sept. 1998.

[4] 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, July 2003.

[5] G. Anastasi, M. Conti, and E. Gregori, “IEEE 802.11 ad hoc networks: protocols, performance and open issues,” in *Ad Hoc Networking*. New York: IEEE Press and John Wiley and Sons, Inc., 2004.

[6] C. L. Fullmer and J. J. Garcia-Luna-Aceves, “Solutions to hidden terminal problems in wireless networks,” in *Proc. of the ACM SIGCOMM*, Cannes, France, Sept. 1997, pp. 39–49.

[7] IEEE Standards Dept., *IEEE Standard 802.11*. IEEE Press, 1999.

[8] M. Takai, J. Martin, and R. Bagrodia, “Effects of wireless physical layer modeling in mobile ad hoc networks,” in *Proc. of the 2nd ACM MOBICOM*, Long Beach, CA, 2001, pp. 87–94.

[9] G. H. Golub and C. F. van Loan, *Matrix Computations*. Baltimore, MD: The Johns Hopkins Univ. Press, 1983.

[10] C. Shen, Y. Zhu, S. Zhou, and J. Jiang, “On the performance of V–BLAST with zero–forcing successive interference cancellation receiver,” in *Proc. of IEEE GLOBECOM 2004*, Dallas, TX, 2004.

[11] M. Levorato, S. Tomasin, P. Casari, and M. Zorzi, “Analysis of spatial multiplexing for cross–layer design of MIMO ad hoc networks,” IEEE VTC 2006 – Spring, Melbourne, Australia, May 2006, 7–10, accepted.

[12] C. A. Balanis, *Antenna Theory: Analysis and Design*, 2nd ed. New York: John Wiley and Sons, Inc., 1996.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE ICC 2006 proceedings.