Physical Layer Approximations for Cross–Layer Performance Analysis in MIMO-BLAST Ad Hoc Networks

Marco Levorato, *Student Member, IEEE,* Stefano Tomasin, *Member, IEEE,*
Paolo Casari, *Student Member, IEEE,* and Michele Zorzi, *Fellow, IEEE*

***Abstract*— In this paper, we consider a MAC protocol for ad hoc networks where nodes are equipped with multiple antennas, and communications are spatially multiplexed using the Bell labs LAyered Space Time (BLAST) system. The contribution of this paper is twofold. First, we introduce two different analytical models aimed at predicting the propagation of detection errors within the BLAST receiver, the first based on a Gaussian approximation of the detection errors, and the second on the weighed enumeration of error configurations. A simplification of the latter, with lower complexity, is also obtained and compared to the original model. We then use these analytical tools to assess the performance of a cross–layer MAC protocol, and compare it with fully detailed simulations. Since the analytical tools replace the simulation of the physical layer, the proposed semianalytical approach is much faster than the bit-by-bit simulations. Numerous results are provided for the network performance assessment, showing that our semianalytical approach is able to predict network behavior with very good accuracy, but much lower complexity.**

***Index Terms*— Ad hoc networks, cross–layer MAC design, MIMO–BLAST, performance analysis.**

I. Introduction

Ad hoc networks are collections of terminals linked together by means of packet radio communications with no underlying infrastructure and no need for supervision by an external network operator. They are fast and cheap to set up, and may represent the only affordable solution in certain scenarios such as battlefields, natural disaster recovery, data sharing during business meetings, etc. [1]. Existing technologies commonly used for ad hoc networks, such as IEEE 802.11a, would enable communications at up to 54 Mbps [2], but in practice this is possible only at very short distances. Moreover, current ad hoc networks are limited by interference among nodes, that is usually prevented by collision avoidance mechanisms. Unawareness of this interference may give rise to problems such as the hidden terminal [3], [4].

Manuscript received April 28, 2006; revised November 9, 2006; accepted February 3, 2007. The associate editor coordinating the review of this paper and approving it for publication was S. Kishore. Part of this work has been presented at IEEE ICC 2006. This work has been supported in part by the U.S. Army Research Office under MURI grant no. W911NF–04–1–0224.

The authors are with the Department of Information Engineering (DEI), University of Padova, Via Gradenigo 6/B, I–35131 Padova, Italy (e-mail: {levorato, stefano.tomasin, casarip, zorzi}@dei.unipd.it). M. Zorzi is also with the University of California at San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0436.

Digital Object Identifier 10.1109/TWC.2007.xxxxx.

The use of multiple antennas for both transmission and reception (multiple input–multiple output, MIMO, systems) is seen as a great opportunity to increase the network throughput and to alleviate interference problems. An important direction of research has been the application of beamforming techniques to ad hoc networks [5], where directional antennas at both the transmitter and the receiver are used to focus the transmitted energy. Among the disadvantages of beamforming, we mention the need to know the channel at the transmitter, with a consequent increase in protocol overhead for the exchange of channel state information among terminals. A second approach is the use of space–time codes to provide spatial diversity. In [6], an Alamouti [7] scheme has been applied to ad hoc networks to increase link reliability. However, spatial diversity translates only partially into a network throughput increase. Spatial multiplexing (SM) instead [8]–[10] increases throughput by transmitting different data on each antenna, and potentially allows for gigabit–per–second wireless links [11]. In this paper we focus on ad hoc networks implementing SM.

In order to assess the network performance, two approaches are usually taken in modeling the physical (PHY) layer. The first approach is to perform a bit-by-bit simulation of each node and the entire network, which obtains accurate results but is very demanding, except for small networks. In the second approach, the physical layer is strongly simplified with various assumptions. For beamforming, antennas' sidelobes are often neglected and the main lobe is modeled with a constant response over a certain azimuthal extension [5], [12]–[14]. For space–time coded systems, terminals are modeled as single-antenna devices and a gain is introduced to account for the increased diversity provided by MIMO. These simplifications may lead to significant inaccuracies [15].

In this paper we consider a MIMO ad hoc network based on SM, in which the receivers implement the Bell labs layered space time (BLAST) technique [8], [9], to cancel interference and improve detection. For the performance assessment, we propose a semianalytical technique based on a closed-form expression of the bit error rate (BER) at a BLAST receiver, for a given channel and network topology. Most existing literature is focused on evaluating the capacity of V-BLAST [16], [17] or the average BER for specific channel statistics [18], [19]. Instead, we consider a general technique for any channel model that provides BER for each detected stream averaged only over noise and data. The analysis includes the effects of error propagation due to erroneous detections, and considers two approaches for BER computation. In the first approach, the residual interference is modeled as additive Gaussian noise (*Gaussian technique*). As observed in [20] and confirmed by our results, this approximation may not be accurate for the BER evaluation of a single link, especially when the transmit antennas are more than the receive antennas of each node. Hence, as a second approach, we consider an interference model based on the exhaustive enumeration of all possible error configurations (*enumeration technique*). In order to reduce complexity, we also propose a variant of this approach, in which we enumerate only a subset of the error configurations, while bounding the error probability for unaccounted events. This last technique is based on pruning the tree describing all possible error configurations (*pruned tree technique*).

The BER statistics derived for the link level are then used to obtain network-level results. We consider a cross–layer physical–medium access control (MAC) protocol to exploit the potential of SM, based on request/clear to send signaling and on a backoff mechanism to limit interference, a preliminary version of which is presented in [21]. We also address the design of the receiver, considering the limited computational capabilities of the nodes. In particular, since BLAST performance is dictated by its ability to detect and cancel interfering streams, it may be beneficial that a receiving node intentionally limits the communications for itself and reserves resources for canceling interfering signals. Hence, we consider various strategies for granting incoming communications and prioritizing signal detections. A comparison of bit-by-bit simulations with our semianalytical approach is provided for these system designs.

Based on our extensive evaluations, we conclude that the Gaussian technique, which may be inaccurate in a single link communication, is indeed the one offering the best trade off between complexity and accuracy when compared to fully simulated network results.

The paper is organized as follows. First we describe the BLAST system in Section II. The performance analysis of BLAST is described in Section III, where we introduce the Gaussian, enumeration and pruned tree techniques. In Section IV we provide a description of the cross–layer PHY–MAC protocol for ad hoc networks based on SM. Extensive numerical results for both link and network level performance are discussed in Section V, comparing bit-by-bit simulations and the semianalytical techniques. Conclusions are outlined in Section VI.

II. System Description

We consider a packet wireless network with nodes having N_A antennas. As transmission scenario we assume a frequency–flat block fading channel. Nodes can both transmit and receive information, in a peer-to-peer fashion. All nodes are in communication range of each other, *i.e.*, we consider a completely connected network. To keep the analysis simple, we assume bit-synchronous node operations, although this is not a requirement when using BLAST receivers (see [10] and references therein).

A. Node operations

Transmitting nodes—When operating as a transmitter, the generic node p splits each packet into sub-packets called *streams*. Each stream is sent from one antenna, so that u_p streams are multiplexed using u_p antennas during transmission. Assuming that $N_{\rm Tx}$ nodes with indices $\{1, 2, \ldots, N_{\rm Tx}\}$ are transmitting, the total number of simultaneously transmitted streams is $U = \sum_{p=1}^{N_{\rm Tx}} u_p$. We can identify each transmit antenna with a transmit antenna index (TAI), starting with the first antenna of the first node and ending with antenna $u_{N_{\rm Tx}}$ of node $N_{\rm Tx}$.

Let us define the column vector $\boldsymbol{s}'(t)=[s'_1(t), \ldots, s'_U(t)]^T$ whose entry $s'_i(t)$ is the symbol transmitted from the antenna with TAI i at time tT, where T is the symbol period and ^{T} denotes the transpose operation.

We assume that node p always transmits with a total power $P_{\rm tot}$ that is uniformly distributed among the u_p antennas. Hence, the power of the transmitted data signal on the antenna with TAI i belonging to node p is $\sigma^2_{s'}(i) = {\rm E}[|s'_i(t)|^2] = P_{\rm tot}/u_p$ where ${\rm E}[\cdot]$ denotes expectation. For the sake of a simpler notation, we omit in the following the time index t in all signals.

Receiving nodes—When operating as a receiver, the generic node q uses all the N_A antennas and the column vector of the N_A received samples can be written as $\boldsymbol{r}^{(q)} = \tilde{\boldsymbol{H}}^{(q)}\boldsymbol{s}' + \boldsymbol{\nu}'^{(q)}$, where $\boldsymbol{\nu}'^{(q)}$ is the column noise vector of length N_A, and $\tilde{\boldsymbol{H}}^{(q)}$ is the $N_A \times U$ channel matrix whose entry $\tilde{\boldsymbol{H}}^{(q)}_{\ell,m}$ represents the complex baseband channel gain between the transmit antenna of TAI m and the ℓth receive antenna.

In order to limit the node complexity, we assume that each receiving node has a partial knowledge of the channels, *i.e.*, node q can only know the channel gains associated to N_d transmit antennas (called *internal antennas*, IA), whose TAIs are in the sub-set $\mathcal{N}^{(q)} = \{n_1, n_2, \ldots, n_{N_d}\}$, for which we assume perfect channel estimation. Without restriction, we assume that the TAIs of known antennas are the first N_d, *i.e.*, $\mathcal{N}^{(q)} = \{1, 2, \ldots, N_d\}$. During data transmission, in general, $\mathcal{N}^{(q)}$ includes the TAI of all granted transmissions intended to reach node q and the TAI of some other interfering transmissions.

We define $\boldsymbol{H}^{(q)}$ as the matrix containing the columns $\tilde{\boldsymbol{H}}^{(q)}_{\cdot,i}$, with $i \in \mathcal{N}^{(q)}$, and we define the N_d-size vector of data symbols belonging to $\mathcal{N}^{(q)}$ as $\boldsymbol{s}^{(q)} = [s'_1, s'_2, \ldots, s'_{N_d}]^T$. The variance of the entries of $\boldsymbol{s}^{(q)}$ is $\sigma^2_{s^{(q)}}(i) = \sigma^2_{s'}(i), \forall i$.

The transmitting antennas whose signals are not detected by node q are instead indicated as *external antennas* (EA) to q. According to $\mathcal{N}^{(q)}$, we assume that the TAI of EA are $N_d+1, N_d+2, \ldots, U$ and the corresponding channel matrix $\bar{\boldsymbol{H}}^{(q)}$ contains the columns $\tilde{\boldsymbol{H}}^{(q)}_{\cdot,i}$, with $i = N_d+1, N_d+2, \ldots, U$. The transmitted symbols of EA to q are denoted as $\bar{\boldsymbol{s}}^{(q)} = [s'_{N_d+1}, s'_{N_d+2}, \ldots, s'_U]^T$. With the definitions of IA and EA, we can rewrite the received signal as

$$\boldsymbol{r}^{(q)} = \boldsymbol{H}^{(q)}\boldsymbol{s}^{(q)} + \bar{\boldsymbol{H}}^{(q)}\bar{\boldsymbol{s}}^{(q)} + \boldsymbol{\nu}'^{(q)} . \tag{1}$$

The explicit modeling of a pure interference term is an important point in our ad hoc scenario. Indeed, terminals may be able to detect only a limited number of signals, or they may decide to neglect low-power interference and reduce processing, *e.g.*, for energy saving purposes.

In order to simplify notation we will omit in the following the index of the receive node $^{(q)}$ in all variables, since we will always refer to a single node.

B. BLAST receiver

In order to extract a sufficient statistics for detection, the receive node multiplies the vector of the received samples by a matrix matched to the channel. By defining the $N_d \times N_d$ matrix $\boldsymbol{R} = \boldsymbol{H}^H \boldsymbol{H}$, where ^{H} is the Hermitian operator, the obtained vector is $\boldsymbol{z} = \boldsymbol{H}^H \boldsymbol{r} = \boldsymbol{R}\boldsymbol{s} + \boldsymbol{\nu} + \boldsymbol{i}_{\mathrm{EA}}$, where $\boldsymbol{\nu} = \boldsymbol{H}^H \boldsymbol{\nu}'$ and $\boldsymbol{i}_{\mathrm{EA}} = \boldsymbol{H}^H \bar{\boldsymbol{H}} \bar{\boldsymbol{s}}$ accounts for the interference due to EA. We recall that the receiver is not required to know the statistics of $\boldsymbol{i}_{\mathrm{EA}}$.

The BLAST receiver performs the detection of the streams in stages. At each stage the stream with the highest signal to noise plus interference ratio (SNIR) is detected, and its contribution is removed from the vector $\boldsymbol{z}$ before the next stage [8], [9]. The ordered TAI set is $\{k_1, k_2, \ldots, k_{N_d}\}$, which is a permutation of the integers $1, 2, \ldots, N_d$.

Let k_i be the TAI of the stream detected at the ith stage and $\boldsymbol{z}(i)$ the vector obtained from $\boldsymbol{z}$ after the removal of the contributions due to streams with TAI in $\mathcal{K}(i) = \{k_1, k_2, \ldots, k_{i-1}\}$ where we set $\mathcal{K}(1) = \varnothing$ and $\boldsymbol{z}(1) = \boldsymbol{z}$. The detection of stream k_i is performed by combining $\boldsymbol{z}(i)$ with the weighing vector $\boldsymbol{w}(i)$ to obtain the sample $\tilde{s}_{k_i} = \boldsymbol{w}(i)^T \boldsymbol{z}(i)$, which is applied to a threshold detector to provide the symbol estimate $\hat{b}_{k_i}$. The estimated symbol is multiplied by the standard deviation of the transmitted symbol to obtain $\hat{s}_{k_i} = \sigma_s(k_i)\hat{b}_{k_i}$. After detection, the contribution of stream k_i is removed from $\boldsymbol{z}(i)$ to obtain

$$\boldsymbol{z}(i+1) = \boldsymbol{z}(i) - \boldsymbol{R}_{\cdot,k_i}\, \hat{s}_{k_i}\,, \quad i = 1, 2, \ldots, N_d - 1\,, \tag{2}$$

where $\boldsymbol{R}_{\cdot,k_i}$ is the k_ith column of $\boldsymbol{R}$.

In the zero forcing (ZF) approach [8], the weighing vector aims at minimizing the interference, regardless of a possible noise enhancement. Let $\boldsymbol{R}(1) = \boldsymbol{R}$ and then compute $\boldsymbol{R}(i)$, $i = 2, 3, \ldots, N_d - 1$ by nulling the k_ith row and column of $\boldsymbol{R}(i-1)$. The weighing vector $\boldsymbol{w}(i)$ is the k_ith column of $\boldsymbol{R}^+(i)$ (the Moore–Penrose pseudoinverse of $\boldsymbol{R}(i)$ [22]), *i.e.*, its mth element is $w_m(i) = [\boldsymbol{R}^+(i)]_{k_i,m}$, $m = 1, 2, \ldots, N_d$.

Provided that the number of receive antennas is larger than the number of residual IA, $N_d - i$, and that $\boldsymbol{R}(i)$ is full rank, the given weighing vector completely cancels the interference due to streams $k_{i+1}, k_{i+2}, \ldots, k_{N_d}$. However, when $N_A < N_d - i$, $\boldsymbol{R}(i)$ has rank smaller than $N_d - i$, and using $\boldsymbol{R}^+(i)$ leaves some residual interference due to IA after weighing. Alternatively, a minimum mean square error (MMSE) criterion could be adopted for the choice of the weighing vector [23], which jointly considers noise and interference. In this paper we only consider the ZF approach, since we verified by simulation that MMSE does not bring any significant advantage and its derivation is entirely analogous to the ZF case.

Lastly, for real constellations, a receiver with improved performance has been derived in [10]. Accordingly, in the forthcoming analysis we shall consider the real part $\mathfrak{Re}\,[\boldsymbol{R}]$ instead of $\boldsymbol{R}$.

III. Performance Analysis

In this section we aim at evaluating the BER of BLAST given the channel matrix $\tilde{\boldsymbol{H}}$. The performance of a SM system is affected by a) imperfect IA cancellation due to detection errors, b) EA interference, c) residual interference due to the pseudo-inverse of $\boldsymbol{R}(i)$, and d) noise. The impact of imperfect cancellation on the detection of forthcoming streams is modeled in this paper with two approaches. In the first approach, denoted *Gaussian technique* and described in Section III-A, the interference is considered as additional Gaussian noise. In the second approach of Section III-B, denoted *enumeration technique*, we exhaustively enumerate all the various configurations of detection errors and compute the conditional BER for each configuration. Although the enumeration technique is more accurate than the Gaussian technique, the number of configurations to be explored increases exponentially with the number of IA, with a consequent increase in computation time. Hence, in Section III-C we derive a suboptimal technique, denoted *pruned tree technique*, where only the most relevant error configurations are explored.

In all cases, EA interference is approximated as Gaussian noise, with autocorrelation matrix

$$\boldsymbol{R}_{\mathrm{EA}} = \mathrm{E}[\boldsymbol{i}_{\mathrm{EA}} \boldsymbol{i}_{\mathrm{EA}}^H] = \boldsymbol{H}^H \bar{\boldsymbol{H}} \bar{\boldsymbol{\Sigma}} [\boldsymbol{H}^H \bar{\boldsymbol{H}}]^H, \tag{3}$$

where $\bar{\boldsymbol{\Sigma}}$ is a diagonal matrix with entries $\bar{\boldsymbol{\Sigma}}_{\ell,\ell} = \sigma^2_{s'}(N_d + \ell)$, $\ell = 1, 2, \ldots, U - N_d$. Then the power of EA interference after weighing can be written as $\sigma^2_{\mathrm{EA}}(k_i) = ||\boldsymbol{w}(i)^T \boldsymbol{R}_{\mathrm{EA}}||^2$.

A. Gaussian technique

We model the residual interference due to errors in the detection process as an *error signal* with Gaussian statistics, zero mean, and variance depending on the error rate. This approach has already been considered in literature, for some particular transmission scenarios. Using the Gaussian approximation the average BER has been obtained for Rayleigh fading channels when $N_d \leq N_A$, [19], [20]. In this section we provide the general expression even for the case of $N_d > N_A$ and for any channel $\boldsymbol{H}$, rather than for a specific channel statistics. This provides a tool for evaluating network performance even when active nodes have different channel statistics.

Assuming a BPSK transmission, the transmitted signal s_k takes a value $\pm\sigma_s(k)$, that depends on the number of active

$$\gamma(k_i) = \frac{\sigma_s^2(k_i)|\boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k_i}|^2}{\frac{N_0}{2}||\boldsymbol{w}(i)^T \boldsymbol{H}^H||^2 + \sigma^2_{\mathrm{EA}}(k_i) \;\; + \displaystyle\sum_{k \in \mathcal{K}(i)} |\boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k}|^2 \sigma_e^2(k) \;\; + \displaystyle\sum_{k \in \mathcal{N} \setminus \mathcal{K}(i+1)} |\boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k}|^2 \sigma_s^2(k)} \tag{6}$$

antennas. Then, the error signal for stream k is

$$e_k = \hat{s}_k - s_k \in \{-2\sigma_s(k), 0, 2\sigma_s(k)\}\,. \quad (4)$$

When the received stream k is affected by a BER $P_e(k)$, $e_k = 0$ with probability $1 - P_e(k)$, $e_k = -2\sigma_s(k)$ with probability $P_e(k)/2$ and $e_k = 2\sigma_s(k)$ with probability $P_e(k)/2$. Note that an erroneous cancellation actually doubles interference. Hence, the variance of the interference signal due to stream k can be written as

$$\sigma_e^2(k) = \mathrm{E}[e_k^2] = 4\sigma_s^2(k)P_e(k) \quad (5)$$

and the SNIR for stream k_i can be written as in (6).

The first term of the denominator in (6) is the power of noise, with power spectral density $N_0/2$ since we are considering a real constellation. The second term of the denominator accounts for the interference due to EA. The power of the residual interference due to imperfect cancellation is provided by the third term in the denominator. With the last term in the denominator we have also inserted the interference due to IA not yet canceled, since when the number of receiving antennas is less than the number of undetected IA, the vector $\boldsymbol{w}(i)$ is not able to completely remove the interference due to streams $k_{i+1}, k_{i+2}, \ldots, k_{N_d}$. Note that, from the definition of $\mathcal{K}(i)$, $\mathcal{N} \setminus \mathcal{K}(i+1)$ is the set of all IA not yet canceled, not including the IA k_i.

The BER for a BLAST transmission over channel $\boldsymbol{H}$ with an uncoded BPSK modulation, is given by $P_e(k_i) = Q\left(\sqrt{\gamma(k_i)}\,\right)$, where $Q(\,\cdot\,)$ is the complementary Gaussian distribution.

B. Enumeration technique

As we will show in Section V-A, the Gaussian technique is not always accurate, especially with an increasing number of transmitting antennas. Hence, we propose here the *enumeration technique* which takes into account the exact interference statistics by enumerating, at stage i of BLAST, all possible error configurations of the previously detected streams $k_1, k_2, \ldots, k_{i-1}$ and their impact on the probability of erroneous detection of stream k_i. A similar approach has been considered in [18] for the computation of the average BER with the assumptions of two transmit antennas, high signal to noise ratio, and Rayleigh fading. Here we generalize the technique for any channel statistics and any number of antennas.

Define the ordered error vector until detection of stream k_i as $\boldsymbol{e}^{(i)} = [e_{k_1}, e_{k_2}, \ldots, e_{k_{i-1}}]^T$, where entries are provided by (4). The error probability on stream k_i can be conditioned on the error configuration of previously canceled streams $\bar{\boldsymbol{e}}$, obtaining, for $i = 2, 3, \ldots N_d$,

$$\mathrm{P}[e_{k_i} = a] = \sum_{\substack{\bar{e}_\ell \in \mathcal{V}_\ell \\ \ell = 1, \ldots, i-1}} \mathrm{P}\left[e_{k_i} = a \middle| \boldsymbol{e}^{(i)} = \bar{\boldsymbol{e}}\right] \mathrm{P}[\boldsymbol{e}^{(i)} = \bar{\boldsymbol{e}}]\,, \quad (7)$$

where $\mathcal{V}_\ell = \{-2\sigma_s(k_\ell), 0, 2\sigma_s(k_\ell)\}$, contains the variance of $\bar{\boldsymbol{e}}_\ell$, $\ell = 1, \ldots, i-1$. Moreover, $a \in \{-2\sigma_s(k_i), 2\sigma_s(k_i)\}$, the summation is taken over all possible error configurations $\bar{\boldsymbol{e}}$ of the $(i-1)$ detected streams, and $\mathrm{P}[\cdot]$ denotes probability. For the first detected stream, we have no previous errors, thus

$$\mathrm{P}[e_{k_1} = +2\sigma_s(k_1)] = \mathrm{P}[e_{k_1} = -2\sigma_s(k_1)] = P_e(k_1)/2\,. \quad (8)$$

The expression (7) can be computed more efficiently by considering the associated tree of error configurations. An example of error tree is shown in Fig. 1 (where for simplicity we set $\sigma_s(i) = 1$, $i = 1, 2, \ldots, N_d$). Starting from the root (level 0), the ith level of the tree corresponds to the detection of stream k_i, and each node corresponds to a different error configuration of the previously detected symbols. In particular, for the first stream k_1, corresponding to the first level after the root, there is only one configuration, since there are no previous detections. Three branches depart from the root, corresponding to the three possible error configurations of $\hat{s}_{k_1}$,

$$\tilde{s}_{k_i}|\boldsymbol{e}^{(i)} = s_{k_i} + \sum_{\ell=1}^{i-1} \boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k_\ell} \boldsymbol{e}_\ell^{(i)} + \sum_{k \in \mathcal{N} \setminus \mathcal{K}(i+1)} \boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k} s_k + \boldsymbol{w}(i)^T \boldsymbol{\nu} + \boldsymbol{w}(i)^T \boldsymbol{i}_{\mathrm{EA}}\,. \quad (9)$$

$$\gamma_{s_{k_i} = \pm\sigma_s(k_i)|\boldsymbol{e}^{(i)}} = \frac{\left| \pm \boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k_i} \sigma_s(k_i) + \sum_{\ell=1}^{i-1} \boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k_\ell} \boldsymbol{e}_\ell^{(i)} \right|^2}{\frac{N_0}{2} ||\boldsymbol{w}(i)^T \boldsymbol{H}^H||^2 + \sigma_{\mathrm{EA}}^2(k_i) + \sum_{k \in \mathcal{N} \setminus \mathcal{K}(i+1)} |\boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k}|^2 \sigma_s^2(k)}\,, \quad (10)$$

$$\mathrm{P}\left[e_{k_i} = 2\sigma_s(k_i) \middle| \boldsymbol{e}^{(i)}, \boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k_i} \sigma_s(k_i) \geq \sum_{\ell=1}^{i-1} \boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k_\ell} \boldsymbol{e}_\ell^{(i)}\right] = Q\left(\sqrt{\gamma_{s_{k_i} = -\sigma_s(k_i)|\boldsymbol{e}^{(i)}}}\right) \quad (11)$$

$$\mathrm{P}\left[e_{k_i} = -2\sigma_s(k_i) \middle| \boldsymbol{e}^{(i)}, \boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k_i} \sigma_s(k_i) \geq -\sum_{\ell=1}^{i-1} \boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k_\ell} \boldsymbol{e}_\ell^{(i)}\right] = Q\left(\sqrt{\gamma_{s_{k_i} = \sigma_s(k_i)|\boldsymbol{e}^{(i)}}}\right) \quad (12)$$

$$\mathrm{P}\left[e_{k_i} = 2\sigma_s(k_i) \middle| \boldsymbol{e}^{(i)}, \boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k_i} \sigma_s(k_i) < \sum_{\ell=1}^{i-1} \boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k_\ell} \boldsymbol{e}_\ell^{(i)}\right] = \left[1 - Q\left(\sqrt{\gamma_{s_{k_i} = -\sigma_s(k_i)|\boldsymbol{e}^{(i)}}}\right)\right] \quad (13)$$

$$\mathrm{P}\left[e_{k_i} = -2\sigma_s(k_i) \middle| \boldsymbol{e}^{(i)}, \boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k_i} \sigma_s(k_i) < -\sum_{\ell=1}^{i-1} \boldsymbol{w}(i)^T \boldsymbol{R}_{\cdot,k_\ell} \boldsymbol{e}_\ell^{(i)}\right] = \left[1 - Q\left(\sqrt{\gamma_{s_{k_i} = +\sigma_s(k_i)|\boldsymbol{e}^{(i)}}}\right)\right] \quad (14)$$

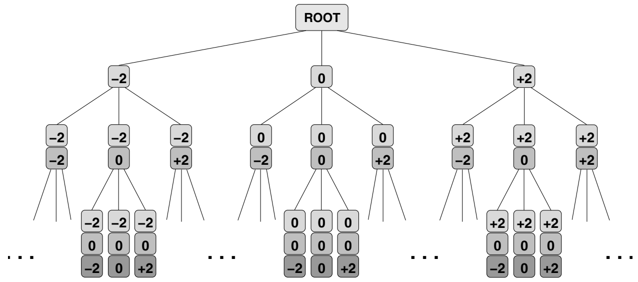

Fig. 1. Error configuration tree for $\sigma_s(i)=1,\ i=1,2,\ldots,N_d$.

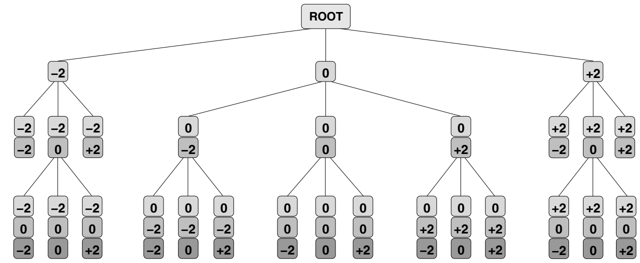

Fig. 2. Pruned error tree for $n_e=2$ and $N_d=4$, $\sigma_s(i)=1,\ i=1,2,3$.

i.e., $-2\sigma_s(k_1), 0, 2\sigma_s(k_1)$. The second level corresponds to the detection of the second stream k_2 and there are three possible error configurations of the previously detected stream k_1.

The general level i has a total of 3^i nodes, each representing one term of the summation in (7). A node at level i is identified by an error configuration vector $\boldsymbol{e}^{(i)}$ and an error value e_{k_i}. For the computation of the error probability of each term of (7) we must consider the signal at the input of the detector, given a specific error configuration. Given an error vector $\boldsymbol{e}^{(i)}$, the signal at the detector input can be written as (9). Let us define the two conditional SNIRs reported in (10), where we observe that the interference due to detection errors produces a shift in the position of the received signal, and hence changes its power in the numerator of the SNIR.

We must now consider the following cases, according to the transmitted signal and the level of interference due to error propagation. The conditional error probability for stream k_i, assuming equally likely data symbols, can be written as in (11)–(14).

The probability of a generic node at level i characterized by the error vector $\bar{\boldsymbol{e}}^{(i)}$ and the error value e_{k_i}, can be obtained as $\mathrm{P}\left[e_{k_i}=\bar{e}_{k_i}|\boldsymbol{e}^{(i)}=\bar{\boldsymbol{e}}^{(i)}\right]$, $i=2,3,\ldots,N_d$, where $\mathrm{P}[e_{k_1}=\bar{e}_{k_1}]$ is provided by (8). This calculation needs averaging over the transmitted symbol, since the same $\boldsymbol{e}^{(i)}$ has a different impact on the error probability, according to the signal sent. Further, averaging $\mathrm{P}\left[e_{k_i}=\bar{e}_{k_i}|\boldsymbol{e}^{(i)}=\bar{\boldsymbol{e}}^{(i)}\right]$ over all $\boldsymbol{e}^{(i)}$ yields the total probability of wrong detection for stream k_i.

Lastly, note that we could also consider the various error configurations of EA for the SNIR evaluation. In this case, since EA are not detected, the error signal would be equal to $-\sigma_{s'}(k)$ or $\sigma_{s'}(k)$ with equal probability $1/2$, for $k=N_d+1, N_d+2,\ldots,U$. Even though this modeling would provide more accurate results than approximating EA as Gaussian noise, it would further increase the computational complexity of the enumeration technique.

C. Pruned tree technique

In order to reduce the computational complexity of the enumeration technique, we limit the exploration of the tree to just a few branches. Sub-trees departing from nodes with more than n_e errors are approximated with an upper bound on the error probability, assuming that when more than n_e errors occur in interference cancellation, the average error probability in forthcoming detections is high. Hence, all nodes of the subtrees having a sub-root with n_e+1 errors are characterized by an error probability of 0.5.

For example, for $n_e=0$ we obtain

$$\begin{aligned}\mathrm{P}[e_{k_i}\neq 0]|_{n_e=0} \quad &\approx \quad \mathrm{P}[e_{k_i}\neq 0|\boldsymbol{e}^{(i)}=\mathbf{0}]\mathrm{P}[\boldsymbol{e}^{(i)}=\mathbf{0}] \\ &+ \quad \frac{1}{2}\left[1-\mathrm{P}[\boldsymbol{e}^{(i)}=\mathbf{0}]\right]. \qquad (15)\end{aligned}$$

This case is similar to the upper bound derived in [24], where it was assumed that given a decision error in an earlier stage of BLAST, the probability of a subsequent decision error is one. In that case the upper bound is as in (15) with 1 instead of 0.5 as weight of the summation. For a general value of n_e, the error configurations with at most n_e errors before the detection of the ith stream are in the set

$$\mathcal{E}^{(i)}(n_e)=\left\{\boldsymbol{e}:\sum_{\ell=1}^{i-1}\frac{|e_{k_\ell}|}{2\sigma_s(k_\ell)}\leq n_e\right\}. \qquad (16)$$

The error probability for node k_i is then approximated as

$$\begin{aligned}\mathrm{P}[e_{k_i}\neq 0]|_{n_e} \quad &\approx \sum_{\bar{\boldsymbol{e}}\,\in\,\mathcal{E}^{(i)}(n_e)}\mathrm{P}[e_{k_i}\neq 0|\boldsymbol{e}^{(i)}=\bar{\boldsymbol{e}}]\mathrm{P}[\boldsymbol{e}^{(i)}=\bar{\boldsymbol{e}}] \\ &+ \quad \frac{1}{2}\left[1-\sum_{\bar{\boldsymbol{e}}\,\in\,\mathcal{E}^{(i)}(n_e)}\mathrm{P}[\boldsymbol{e}^{(i)}=\bar{\boldsymbol{e}}]\right]. \qquad (17)\end{aligned}$$

The last term weighed by $1/2$ accounts for the probability of configurations with more than n_e errors. We observe that the error propagation tree can then be pruned of all nodes that have more than n_e errors, since subtrees departing from these nodes are approximated with an error probability of 0.5. Fig. 2 shows an example of a pruned tree for $n_e=2$, and $N_d=4$ total IA streams, where at most one detection error is explicitly accounted for and any configuration with more than one error is assumed to yield a correct detection probability of 0.5 in all subsequent stages (corresponding to dead leaves in the tree).

D. Computational complexity comparison

Various algorithms have been proposed for an efficient implementation of BLAST. In [25] a recursive algorithm for the matrix inversion has been proposed. In [26] a square-root based algorithm is proposed, which has been further optimized in the improved square-root (ISR) algorithm [27]. Here we consider ISR as the most efficient and reliable technique, and hence all simulations and semianalytical techniques have a common base complexity of $(2/3)N_d^3+(7/2)N_d^2N_A+\mathcal{O}(N_d^2+N_dN_A)$ complex multiplications.

For the Gaussian technique, the function Q is computed N_d times, to obtain the corresponding BERs. The computation of Q can be performed by a look-up table and we bound the complexity of each function computation with the equivalent of one complex multiplication. Moreover, SNIR must be computed N_d times, according to (6), which requires

$$\begin{aligned} C_\gamma &= (N_A+3+1) + [(1+1+N_A) + N_d(N_A+1) + 1] \\ &= N_d N_A + \mathcal{O}(N_d+N_A) \,. \end{aligned} \tag{18}$$

products and ratios. Hence, the Gaussian technique has a complexity

$$C_G = (C_\gamma+1)N_d = N_d^2 N_A + \mathcal{O}(N_d^2+N_A N_d) \,. \tag{19}$$

For the enumeration technique, we must compute N_d BER functions for each leaf of the error tree. The computation of the conditional SNIRs (10) has the same complexity as (6). The complete error tree has 3^{N_d-1} leaves, leading to an overall complexity

$$C_E = 2 \cdot 3^{N_d-1} C_G \,. \tag{20}$$

For the pruned tree technique, the number of explored leaves is $\lambda(n_e) = \sum_{k=0}^{n_e} \binom{N_d}{k} 2^k$. Hence, the complexity of the pruned error tree algorithm is

$$C_P(n_e) = 2 \cdot \lambda(n_e) C_G \,. \tag{21}$$

Note that in (20) and (21) the factor 2 accounts for the double computation of Q as required by (11)–(14). We observe that the complexity of the enumeration and pruned tree techniques grows exponentially with N_d, while the Gaussian technique grows as its third power. For example, with $N_d = 8$ the ratio between the enumeration and the Gaussian technique is $C_E/C_G = 39366$, while for $n_e = 1$ we have $C_P(1)/C_G = 42$. We conclude that the enumeration technique has a significantly higher complexity than the Gaussian technique and it may be even more demanding than the bit-by-bit simulation. The pruned tree technique instead involves a limited increase of complexity with respect to the Gaussian technique.

IV. Cross–layer Design and Evaluation of a MIMO Ad Hoc Network

In this Section, we address the problem of designing an ad hoc network based on SM. Besides describing protocols and network operations, we highlight the importance of knowing the behavior of the underlying physical layer for a proper MAC design. To this extent, we first illustrate the general network operation, and then give details on cross–layer interactions that improve the overall system performance.

A. Network architecture

A well–designed MAC protocol would exploit the advantages of SM, namely i) an increased number of supported transmissions, *i.e.*, higher spatial reuse, because many links may coexist in the same neighborhood; ii) the ability to send independent data streams through more than one antenna, thus increasing the raw bit rate available between nodes; and iii) the ability of each transmitter to address more receivers simultaneously, using disjoint subsets of antennas for different receivers.

We describe a MAC design that exploits the potential of SM while keeping interference under control. Our protocol uses a handshake procedure to initiate data transmissions and a backoff strategy to regulate the persistency of a node in accessing the channel.

In the following, we will assume that each transmission is preceded by a training sequence and that all nodes know all training sequences used by any neighbor. Different neighbors use different training sequences. Still, due to technological limitations, each node can estimate the channel conditions of up to N_d^{MAX} incoming signals.

B. Handshake procedure

Before data transmissions, request-to-send (RTS) and clear-to-send (CTS) messages are exchanged between the transmitting and receiving nodes. This negotiation phase, besides providing request notifications and transmission grants, also allows each receiver to estimate the local traffic intensity. After the negotiation phase, data streams are sent by each transmitter that obtained a CTS, and finally acknowledgement (ACK) messages are used to confirm the correct packet receptions.

We assume that transmissions are organized in frames, defined as the succession of an RTS, a CTS, a DATA and an ACK, all of fixed predefined length. Nodes are frame–synchronous, so that all transmitters send packets simultaneously in the assigned frame portion as prescribed by the aforementioned structure.

Even though frame–level synchronization may seem a drawback of our approach, we note that i) slot synchronization is currently achieved in protocols belonging to the 802.11 family [2]; ii) for a completely connected network, synchronization is easier to achieve; iii) the constraint can be relaxed if we consider unsynchronized transmissions as EA. Studying the impact of unsynchronized transmissions on the protocol performance is beyond the scope of this paper, and is currently being investigated. Among the advantages of sending all RTSs simultaneously we mention that the receivers are able to collect all requests and form CTSs accordingly. We describe hereon the rules RTSs and CTSs are subject to.

RTS policy—Each RTS message may contain multiple requests, each specifying the number of streams that would be sent toward their own destinations. This allows transmitters to require link set up with different recipients in a single frame. Nonetheless, we remark that the higher bit rate achieved through a higher number of multiplexed streams has two drawbacks: i) the available transmit power is to be split among all used antennas, resulting in lower SNIRs at the receiver; ii) detection performance in the presence of a high receiver load decreases very quickly with SNIR [21]. When trading off load (*i.e.*, number of antennas) with communication reliability, any transmitter has to take the described detrimental effects into account, and make sure that the total number of transmit antennas is suitably chosen.

CTS policy—Upon detecting RTSs, each receiver generates a single CTS message for all the granted requests, based on received RTS content. The criterion for the CTS generation aims at guaranteeing that i) most streams are granted for *network throughput* purposes, and ii) the granted streams are sufficiently protected from interference, so that they can be correctly detected with high probability. In order to limit interference, we apply a so–called *Follow Traffic* (FT) CTS policy [21]. Recall that the receiver has to account for channel estimation resources that will be needed for all detected streams. At first, node D sorts all requests contained in every correctly decoded RTS in order of decreasing received power, regardless of the destination node. The first requests accounting for up to N_d^{MAX} streams are considered for detection and all RTSs for D are granted within this set. In the special case when all requests to D are not within the first N_d^{MAX} requests with the highest power, the FT policy forcedly grants the request with the highest power directed to D and allocates the remaining resources to detect the streams with the highest power not directed to D. This policy strives to guarantee some minimum throughput by granting at least one transmission, while protecting from strong interference. Note that this selection criterion has the advantage of being distributed, and does not require any further signaling among nodes. On the other hand, a centralized selection could achieve other objectives, such as throughput maximization, and would account for the effective SNIR determined through the knowledge of the actual transmissions, rather than of the requests only. However, a centralized algorithm is in contrast with the principle of ad hoc networks and would require some expensive signaling. We therefore consider here only the distributed selection criterion.

Finally, it is worth highlighting that FT relies on power measurements for each stream from the physical layer, so that the MAC layer can take decisions about the streams to be detected or canceled as interference. These decisions are then fed back to the physical layer, which spatially demultiplexes incoming streams based on MAC layer directives. Because of these interactions, FT is a very good example of cross–layer network operation.

C. Backoff strategies

Even in the presence of a conservative policy for protecting data from superimposed streams, transmissions may indeed fail, *e.g.*, due to channel impairments or excessive interference. In order to prevent nodes from persisting in transmitting toward congested or unavailable nodes, in this paper we resort to backoff. In particular, we distinguish between exponential *destination–wise* and *transmitter–wise* backoff. In the destination–wise version, for each request whose CTS is not received, the transmitter defers communication with *that* destination for a random number of frames, randomly distributed in the back-off window $[1, B_{max}]$; in the transmitter–wise version, *all* communications originating from the failing node are deferred regardless of their recipient, as usual in many backoff schemes, *e.g.*, 802.11 DCF [2]. The backoff window is exponentially increased, following the relation $B_{max} = W \cdot 2^{N_f - 1}$ where N_f is the number of subsequently failed attempts. If a maximum preset value N_f^{MAX} is reached, the backoff window is no longer increased. Moreover, packets are assigned a maximum number of attempts, after which they are dropped. The choice of N_f^{MAX} and B_{max} is meant to ensure that a certain number of backoffs with window length B_{max} take place before packet discarding. In the following, we shall refer to destination– and transmitter–wise backoff as *dest–lock*, and *node–lock*, respectively.

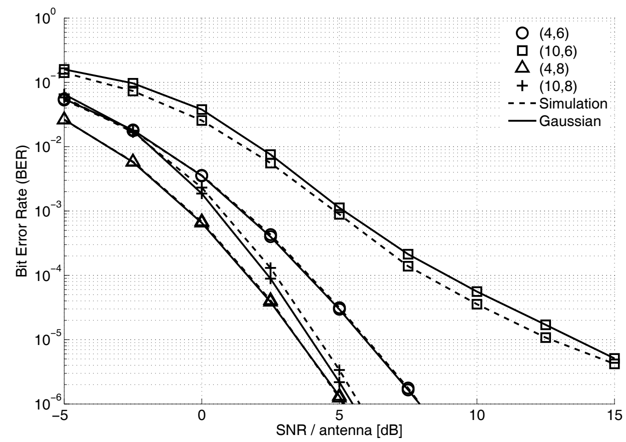

Fig. 3. Comparison between analytical (using the Gaussian technique) and simulated BER results for various configurations (N_d, N_A), with N_d the detected streams (all assumed to be IA) and N_A the receive antennas.

V. Numerical Results

In this Section, we carry out relevant physical level as well as network level simulations that exploit the analytical framework deployed in Section III, and compare results with fully detailed bit–level simulations. A very important conclusion that can be drawn from our results is that a complete and accurate physical layer modeling is crucial for both deciding how different layers interact and tuning relevant communication parameters, since oversimplified models may hamper the statistical significance of simulation results. PHY results are reported in Section V-A, whereas MAC level metrics are given and commented in Section V-C.

A. Bit error rate

In this section we compare the simulated performance of a generic node in the network with the analytical results derived in Section III, considering channel gains having complex Gaussian statistics with zero mean and unit variance. We have compared the simulated results for different configurations of transmitted streams N_d and receive antennas N_A. The receiving node is assumed to detect all streams so that N_d is also the number of detected streams. We begin by examining this simpler setup as it corresponds to a special-case network with lower load, where all transmitters are placed at the same distance from the receiver, and each of them uses a single antenna at the full available power. Even though this setting may not be representative of a real network, it is a significant benchmark to consider before moving to more realistic scenarios.

In Fig. 3 we plot results of simulation and Gaussian technique for the case of a single link. The BER is averaged over all detected streams and various channel realizations, and is shown as a function of the average signal to noise ratio (SNR) per antenna. The average BER obtained with the pruned tree technique is shown in Fig. 4.

Results predicted by the Gaussian approximation technique differ from the simulated ones for two main reasons: $i)$ the interference due to error propagation is not Gaussian, due to different received power levels and signal ordering, and

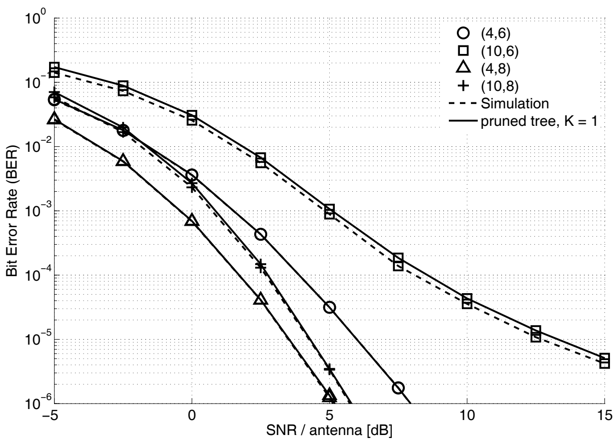

Fig. 4. Performance of the pruned tree technique for BLAST BER evaluation for various configurations (N_d, N_A), with N_d the detected streams (all assumed to be IA) and N_A the receive antennas.

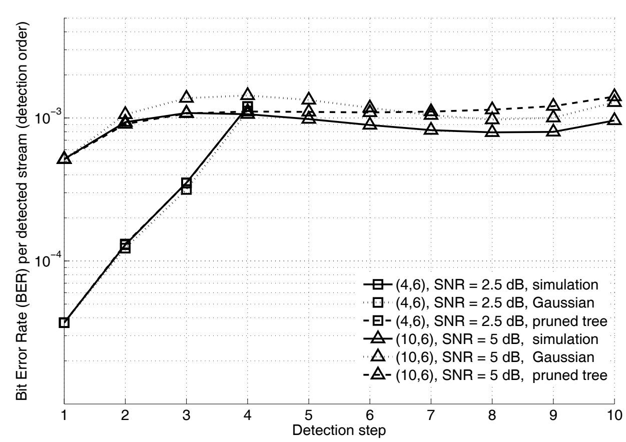

Fig. 5. Analytical and simulated BER *per detected stream* for antenna configurations $(4, 6)$ and $(10, 6)$ and $\mathrm{SNR} = 2.5$ dB and 5 dB, respectively. The abscissa lists the stream index as per the detection order.

$ii)$ the interference due to imperfect ZF is also not Gaussian. Therefore, the approximation is more accurate at higher SNR, where fewer errors occur and the error propagation phenomenon is reduced [20]. Moreover, as long as there are fewer streams than receive antennas, ZF is effective in removing all interference and analysis is in accordance with simulation. When $N_d > N_A$, ZF BLAST yields additional interference that is not Gaussian, providing a further mismatch with simulated results.

The evaluation of the average BER, not reported in a separate figure for conciseness, confirms that the enumeration technique has a better match than the pruned error tree approach. For the enumeration technique, simulated and analytical BER exhibit a perfect match when $N_d \leq N_A$, while for $N_d > N_A$ a slight mismatch is present, due to the approximation of imperfect ZF as Gaussian interference.

Results reported until now are for the average BER over all detected streams. However, the BER of individual streams may differ from the average behavior, due to different level of interference and power. The per-user BER, averaged with respect to the channel statistics, is shown in Fig. 5 for two antenna configurations. The stream index refers to the detection order and variations in the BER among different users result from the combination of different received powers, error propagation and mutual interference. It is important to observe that the Gaussian approximation can be more accurate than the pruned error tree for the estimate of the BER of some streams. Note in fact that approximating the error probability as 0.5 after n_e errors flattens the analytical BER curve of the pruned tree approach.

B. Network simulation environment

The behavior of a network in a general arrangement can not be directly derived from the BER results given in the previous subsection. Moreover, network performance directly depends on MAC choices. Hence, BER results can not provide the full picture about the accuracy of the analytical methods, since relevant differences at the physical layer may be smoothed out or amplified by network behaviors. Before designing or optimizing protocols based on pseudo-analytical results, it is therefore very important to assess the validity of the approximate techniques in a more general networking scenario.

To this end, we arrange a total of 25 nodes on a grid in a $(100 \times 100)\ \mathrm{m}^2$ square area, such that the distance between nearest neighbors is $25\,\mathrm{m}$. With this setting, we ensure that the error probability of a transmission between the two nodes at the largest possible distance is below 1%, in the absence of interference. Assuming transmission in the 5.8 GHz band, we can obtain independent fading by packing 8 antennas per node within nearly $20\,\mathrm{cm}$, which fits on the screen of a laptop computer. Considering indoor transmission, we assumed a path loss attenuation coefficient of 4.

Packets are generated according to a Poisson process of rate λ packets per second per node. Each packet is randomly assigned a length of $k \times 1000$ bits, with k uniformly chosen in the set $\{1, 2, 3, 4\}$. Nodes keep backlogged packets in a queue that can store up to 120 1000-bit units. Nodes can track at most $N_d^{MAX} = 32$ training sequences. Transmissions are framed as explained in Section IV, and all nodes share the same frame synchronization. Since the frame structure is fixed, the length of signaling packets and data streams is also fixed to be 200 bits for RTSs, CTSs and ACKs and 1000 bits per used transmit antenna for SM data streams, including PHY preambles used, *e.g.*, for channel estimation. The overall frame duration is then 1600 bits, comprising RTS, CTS, data and ACK. Under the preceding assumptions, channels may be assumed to be constant over the whole duration of a frame. Recall that, in order to save overhead, RTSs and CTSs may contain transmission requests and grants to multiple nodes, respectively. Moreover, terminals confirm each stream individually, so that correctly received streams are acknowledged even if other streams belonging to the same packet are not received correctly.

All signaling packets are sent with a single antenna, since in this phase we are not interested in maximum throughput, but rather in maximum probability of success. Sending short packets with one antenna allows both to increase the transmission power and to reduce network load, which results in increased reliability. In this case, RTSs, CTSs and ACKs can travel long distances without errors, maximizing also the

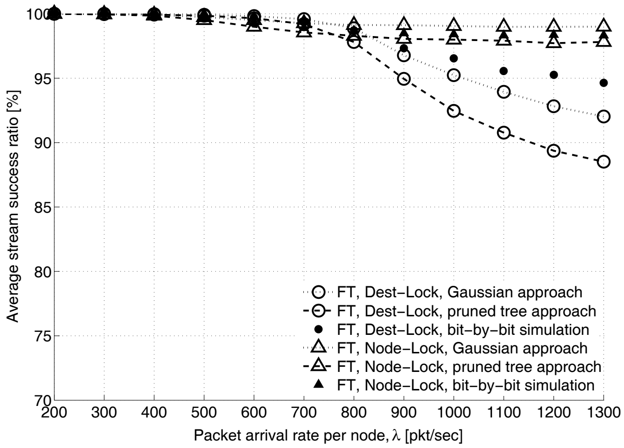

Fig. 6. Average 1000-bit stream transmit success ratio as a function of λ, dest–lock and node–lock.

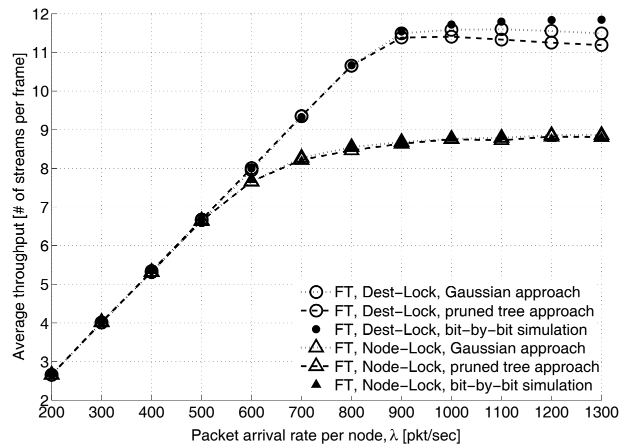

Fig. 7. Average network throughput as a function of λ, dest–lock and node–lock.

nodes' awareness about neighboring activity, according to the value of N_d^{MAX}. We use the distributed ad hoc MAC protocol described in Section IV to obtain network simulations, where signaling among nodes is limited to RTSs and CTSs and no other information is assumed.

C. Network results

A key point of our analysis is to verify whether and to what extent using a given approximation method corresponds to accurate networking results, since MAC protocols are likely to have a smoothing effect on discrepancies arising from lower level models. In order to do so, bit-by-bit simulations of the whole system, including complete physical and MAC layer modeling, are provided and compared with those obtained using the Gaussian and the pruned tree approximations. The complete tree exploration is not considered for network results, since it has an exceedingly high complexity. Note that the pruned tree approach provides a better match to simulated average BER than the Gaussian approximation technique, while the Gaussian approximation better follows BER behavior for individual streams. On the other hand, the Gaussian approximation is much less complex than the pruned tree approach.

The BER results for the single link are closely related to the PER and its complementary function, the network average success rate (1 − PER). In Fig. 6 we compare the average success rate obtained by simulation and analytical techniques, considering the average ratio between the number of correctly detected 1000-bit streams and the total number of transmitted streams per frame. We observe that the conclusions derived from the average BER results for a single link are modified by the MAC behavior and the accuracy of the analysis for each stream becomes more relevant. In particular, the pruned tree technique, which provides a good BER match, yields PER values that differ significantly from simulation results due to the limited number of branches considered. On the contrary, the Gaussian technique performs well under both dest–lock and node–lock policies and under all considered traffic levels, since as the number of streams increases, the distribution of interference becomes Gaussian by the central limit theorem.

In order to explore the accuracy of the analytical techniques we considered other network metrics, *i.e.*, average network throughput, average queue length, and delay.

Fig. 7 shows the average network throughput, defined as the number of correctly detected 1000-bit streams per frame, as a function of the offered traffic λ. As in Fig. 6, the analytical techniques are accurate for low traffic, and the pruned error tree approach exhibits a mismatch with respect to simulated results due to flattening of the BER per user.

Figs. 8 and 9 depict the average queue length and the delay, defined as the average time elapsed from the packet generation to the ACK reception following a successful packet transmission. We note that also in this case network parameters are well approximated by the MAC protocol evaluated by the Gaussian technique. About the latency, the smaller throughput predicted by the pruned tree translates into a greater delay in the dest–lock case. On the contrary, with node–lock, both approaches provide very close approximations. Fig. 10 shows the average number of transmit/receive data links that a single node activates per frame, possibly containing more than one 1000-bit stream transmission per link. Both the pruned tree and the Gaussian techniques are indeed very accurate in the dest–lock and the node–lock cases.

We conclude that the Gaussian approximation is well suited to predict network behavior for a wide range of traffic intensities and different MAC policies, and also has the advantage of a limited complexity when compared with the pruned error tree approach. It is then suitable to obtain fast results, for example to compare the two proposed MAC polices. We have also performed simulations in order to assess the dependence of throughput on the number of antennas at each node. Results, not reported here for conciseness, indicate that throughput is roughly linearly dependent on the number of antennas. An in-depth design of MAC policies can be found in [28]: here we highlight the main differences of the two approaches. The dest–lock policy favors transmissions, as it blocks communications only toward a single unavailable receiver each time a failure occurs. As more nodes transmit simultaneously, receivers must perform more cancellation stages, which results in a greater probability of detection errors. This explains the slight throughput decrease of the dest–lock policy. On the

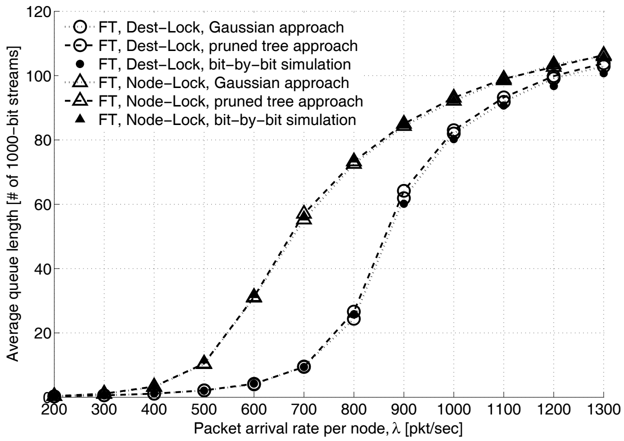

Fig. 8. Average queue length as a function of λ, dest–lock and node–lock.

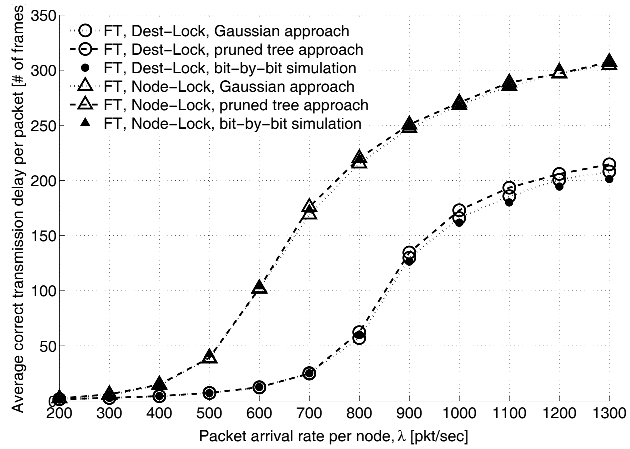

Fig. 9. Average correct transmission delay as a function of λ, dest–lock and node–lock.

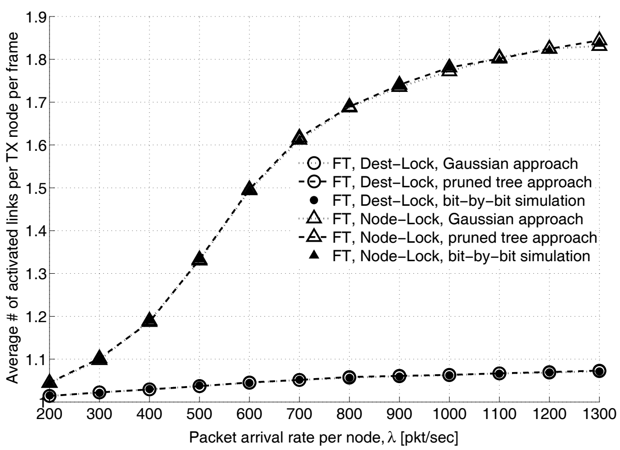

Fig. 10. Average number of activated links per transmitting node per frame as a function of λ, dest–lock and node–lock.

other hand, from Fig. 10, dest–lock generates more RTSs than node–lock, and for very high packet arrival rates dest–lock may cause more collisions than node–lock. In fact, while node–lock is conservative in backing off the entire node, dest–lock better exploits the available links, letting more nodes transmit, each with a lower number of active links.

VI. Conclusions

We have studied analytical approximation techniques for fast performance evaluation of ad hoc networks using BLAST. The match of the approximations with the results obtained by detailed simulations has been assessed both for a single link and for network performance. We have also considered a cross–layer protocol that, by letting the physical and MAC layers interact, can effectively control the interference of an ad hoc network based on spatial multiplexing. From the numerical results, and considering the dramatic decrease of complexity achieved using the approximation instead of bit-by-bit simulations, we can conclude that these approximations provide a useful tool for the design and evaluation of future ad hoc networks using multiple antennas. Indeed, the Gaussian approximation of the interference shows the best match with network simulation results and is to be preferred to error enumeration techniques, that on the other hand show better results for a single link.

Moreover, the considered cross–layer protocol design exhibits promising performance in exploiting spatial multiplexing and provides a significant network throughput increase over existing MAC protocols for the same physical layer.

References

[1] C. Silva Ram Murthy and B. S. Manoj, *Ad Hoc Wireless Networks–Architectures and Protocols*. Upper Saddle River, NJ: Prentice Hall, June 2004.

[2] *ANSI/IEEE Standard 802.11*, IEEE Press, 1999.

[3] 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.

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

[5] R. Ramanathan, “On the performance of ad hoc networks with beamforming antennas,” in *Proc. ACM MOBIHOC*, Oct. 2001, pp. 95–105.

[6] M. Hu and J. Zhang, “MIMO ad hoc networks: Medium access control, saturation throughput, and optimal hop distance,” *J. Commun. Networks*, pp. 317–330, Dec. 2004.

[7] S. M. Alamouti, “A simple transmit diversity technique for wireless communications,” *IEEE J. Select. Areas Commun.*, vol. 16, no. 8, pp. 1451–1458, Oct. 1998.

[8] G. J. Foschini, “Layered space-time architecture for wireless communication in a fading environment when using multiple antennas,” *Bell Labs Tech. J.*, vol. 1, no. 2, pp. 41–59, 1996.

[9] 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. IEEE ISSSE*, Sep. 1998, pp. 295–300.

[10] 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.

[11] 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.

[12] A. Nasipuri, S. Ye, J. You, and R. E. Hiromoto, “A MAC protocol for mobile ad hoc networks using directional antennas,” in *Proc. IEEE WCNC*, Sep. 2000, vol. 2, pp. 1214–1219.

[13] R. R. Choudhury, X. Yang, R. Ramanathan, and N. H. Vaidya, “Using directional antennas for medium access control in ad hoc networks,” in *Proc. ACM MOBICOM*, Sep. 2002, pp. 59–70.

[14] T. Korakis, G. Jakllari, and L. Tassiulas, “A MAC protocol for full exploitation of directional antennas in ad hoc wireless networks,” in *Proc. ACM MOBIHOC*, June 2003, pp. 98–107.

[15] M. Takai, J. Martin, and R. Bagrodia, “Effects of wireless physical layer modeling in mobile ad hoc networks,” in *Proc. ACM MOBIHOC*, Oct. 2001, pp. 87–94.

[16] A. Gorokhov, D. A. Gore, and A. J. Paulraj, "Receive antenna selection for MIMO spatial multiplexing: theory and algorithms," *IEEE Trans. Signal Processing*, vol. 51, no. 11, pp. 2796–2807, Nov. 2003.

[17] *Space-Time Processing for MIMO Communications*. Chichester, England (UK): John Wiley and Sons, Inc., 2005, A. B. Gershman and N. D. Sidiropoulos, *eds*.

[18] S. Loyka and F. Gagnon, "Performance analysis of the V-BLAST algorithm: an analytical approach," *IEEE Trans. Wireless Commun.*, vol. 3, no. 4, pp. 1326–1337, July 2004.

[19] C. Shen, Y. Zhu, S. Zhou, and J. Jiang, "On the performance of V-BLAST with zero-forcing successive interference cancellation receiver," in *Proc. GLOBECOM*, Nov. 2004, pp. 2818–2822.

[20] K. Liu and A. M. Sayeed, "An iterative extension of BLAST decoding algorithm for layered space-time signals," *IEEE Trans. Commun.*, vol. 53, no. 10, pp. 1754–1761, Oct. 2005.

[21] P. Casari, M. Levorato, and M. Zorzi, "On the implications of layered space–time multiuser detection on the design of MAC protocols for ad hoc networks," in *Proc. IEEE PIMRC*, Sep. 2005, pp. 1354–1360.

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

[23] G. Ginis and J. M. Cioffi, "On the relation between BLAST and the GDFE," *IEEE Commun. Lett.*, vol. 5, no. 9, pp. 364–366, Sep. 2001.

[24] R. Narasimhan, "Error propagation analysis of V–BLAST with channel estimation errors," *IEEE Trans. Commun.*, vol. 53, no. 1, pp. 27–31, Jan. 2005.

[25] J. Benesty, Y. Huang, and J. Chen, "A fast recursive algorithm for optimum sequential signal detection in a BLAST system," *IEEE Trans. Signal Processing*, vol. 1, no. 1, pp. 1722–1730, July 2003.

[26] B. Hassibi, "An efficient square-root algorithm for BLAST," in *Proc. IEEE ICASSP*, June 2000, pp. 737–740.

[27] H. Zhu, Z. Lei, and F. P. S. Chin, "An improved square-root algorithm for BLAST," *IEEE Signal Processing Lett.*, vol. 11, no. 9, pp. 772–775, Sep. 2004.

[28] M. Levorato, S. Tomasin, P. Casari, and M. Zorzi, "Analysis of spatial multiplexing for cross–layer design of MIMO ad hoc networks," in *Proc. IEEE VTC-Spring*, May 2006, pp. 1146–1150.

Marco Levorato (S'06) was born in Venice on March 18th, 1980. He obtained both the BE (Electronics and Telecommunications Engineering) and the ME (Telecommunications Engineering) *summa cum laude* from the University of Ferrara (Italy) in 2002 and 2005, respectively. During 2005 he held a fellowship at the University of Padova (Italy), and from January 2006 he has been a Ph.D. student in Information Engineering at the University of Padova under the supervision of Prof. Michele Zorzi. His research interests include cooperative communications, the design of ad hoc networks with multiuser detection, and the analysis of hybrid ARQ techniques.

Stefano Tomasin (S'99-M'03) received the Laurea degree and the Ph.D. degree in Telecommunications Engineering from the University of Padova, Italy, in 1999 and 2002, respectively. In the academic year 1999–2000 he was on leave at the IBM Research Laboratory, Zurich, Switzerland, doing research on signal processing for magnetic recording systems. In the academic year 2001–2002 he was on leave at Philips Research, Eindhoven, the Netherlands, studying multicarrier transmission for mobile applications. He joined the University of Padova first as a contractor researcher for a national research project (2002) and then as an assistant professor (2005). In the second half of 2004 he was visiting faculty at Qualcomm, San Diego (CA) doing research on receiver design for mobile cellular systems. His current research interests include signal processing for wireless communications, access technologies for multiuser/multiantenna systems, and cross–layer protocol design and evaluation.

Paolo Casari (S'05) was born in Ferrara, Italy, on August 20th, 1980. He received the Laurea degree (BE) in Electronics and Telecommunications Engineering (2002) and the Laurea Specialistica degree (ME) in Telecommunications Engineering (2004) *summa cum laude*, both from the University of Ferrara. From September to December 2004, he was with the same university, studying geographic protocols for wireless sensor networks. In January 2005, he joined the Department of Information Engineering at the University of Padova, Italy, where he is currently a Ph.D. student under the supervision of Prof. Michele Zorzi. He held teaching assistantships at both universities. His main research interest is cross–layer protocol design for wireless networks through PHY/MAC/routing interactions, with a particular focus on MIMO ad hoc networks and wireless sensor networks. He is currently on leave at the Massachusetts Institute of Technology, Cambridge, MA, working on energy–efficient protocol design for underwater acoustic sensor networks.

Michele Zorzi (S'89-M'95-SM'98-F'07) was born in Venice, Italy, in 1966. He received the Laurea degree and the Ph.D. degree in electrical engineering from the University of Padova, Italy, in 1990 and 1994, respectively. During the academic year 1992/93, he was on leave at the University of California, San Diego (UCSD), attending graduate courses and doing research on multiple access in mobile radio networks. In 1993, he joined the faculty of the Dipartimento di Elettronica e Informazione, Politecnico di Milano, Italy. After spending three years with the Center for Wireless Communications at UCSD, in 1998 he joined the School of Engineering of the University of Ferrara, Italy, and in 2003 he joined the Department of Information Engineering of the University of Padova, Italy, where he is currently a professor. His present research interests include performance evaluation in mobile communications systems, random access in mobile radio networks, ad hoc and sensor networks, and energy constrained communications protocols.

Dr. Zorzi was the Editor-In-Chief of *IEEE Wireless Communications Magazine* from 2003 to 2005, and currently serves on the steering committee of the IEEE TRANSACTIONS ON MOBILE COMPUTING, and on the editorial boards of the IEEE TRANSACTIONS ON COMMUNICATIONS, the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS, the *Wiley Journal of Wireless Communications and Mobile Computing* and the *ACM/URSI/Kluwer Journal of Wireless Networks*. He was also guest editor for special issues in *IEEE Personal Communications Magazine* (Energy Management in Personal Communications Systems) and the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS (Multi-media Network Radios).