Impact of Space-Time Block Codes on 802.11 Network Throughput

Anastasios Stamoulis and *Naofal Al-Dhahir*

AT&T Shannon Laboratory, Florham Park, NJ

{as,naofal}@research.att.com

Contact Author is Naofal Al-Dhahir, AT&T Shannon Lab, Bldg. 103, Rm C289, 180 Park Avenue, Florham Park, NJ 07932. T:(973)360-8407. F:(973)360-8178.

Abstract

Space-Time Block Codes (STBC) represent a recent, exciting development in the field of physical layer (PHY) data transmission. By employing more than one antenna at the transmitter, and by properly coding data across the transmit antennas, STBC-equipped PHYs promise increased data rates with minimal decoding complexity at the receiver. This article presents a comprehensive study of how the STBC gains at the PHY translate to significant network performance improvement in 802.11a Wireless LANs. We base our study on a detailed, across-all-layers, simulation of an 802.11a system; we have extended the `ns` simulator with an implementation of the 802.11a PHY, which allows us to assess the impact of STBC not only at the PHY layer, but at the higher layers as well. An extensive set of simulations illustrates the merits of STBC, and sheds light on how performance can be improved for TCP traffic. Essentially, STBC presents to TCP a "smoother" wireless channel; the latter phenomenon is corroborated by a brief theoretical analysis as well.

Keywords

space-time block codes, TCP, ns simulator, 802.11, Wireless LANs

I. Introduction

Recent advances in wireless communications promise increased data rates for Wireless LANs. Architectures such as the IEEE 802.11a specify data rates as high as 54Mbit/s, and lay the foundation for an entire new suite of applications. For example, there is ongoing research activity on delivery of high-quality multimedia in a home-networking environment, advanced collaboration tools in office environments, and, ultimately, the creation of a wireless link as fast and reliable as an early-90's wired link.

Undoubtedly, there are a lot of research issues that need to be addressed before such an ultimate goal is realized. A huge challenge is imposed by the adverse nature of the wireless channel. Wireless channels are susceptible to frequency selectivity (which arises from multipath propagation), time selectivity (which arises from movement of the communicating nodes), and exhibit bit error rates (BERs) orders of magnitude higher than those in wired links. The wireless channel impairments not only affect PHY performance, but their effect reverberates across all higher layers. For example, TCP performance deteriorates in the event of lost packets, and fair-queuing scheduling algorithms need to be modified in wireless set-ups (see, e.g., [1]).

Indeed there have been intense research efforts to mitigate the deleterious effects of the wireless channel (e.g., by developing proxy agents). However, intuitively, it seems that instead of working around the defects of the wireless channel, an alternative more effective approach would be to make the wireless channel appear "smoother" to higher layers. This is exactly one of the goals of transmit/receive PHY diversity, which builds upon the apparently intuitive fact that

> Instead of coping with only one "bad" wireless channel, perhaps it is better to create many "bad" wireless channels, and then combine these "bad" wireless links so as to make out of them one "good" wireless channel.

Behind this simple statement, there is more than meets the eye, and there is an extensive body of pertinent work in the Communications–Information Theory literature. From a high-level viewpoint, each bad wireless channel of the aforementioned quote corresponds to a

(transmit-antenna, receive-antenna)

pair. For example, if the transmitter node has two antennas, then the data packets travel to the receiver over two different channels – if they are properly combined, the probability that data packets get lost becomes much smaller.

A special form of transmit/receive PHY diversity are Space Time Block Codes (STBC). STBC combine elegantly techniques from Coding Theory, Matrix Algebra, and Signal Processing to yield transmit/receive diversity gains with minimum computational complexity [2]. Since their invention, STBC have sparked a wide interest in the Communications–Information Theory community, because they promise a significant PHY performance improvement which can potentially increase the data rates in wireless communications [3].

In this work, we set out to investigate the impact of STBC on 802.11a Wireless LANs, and along the way we make two important contributions

- First, we simulate the 802.11a PHY and compare it to an "STBC-enhanced" 802.11a PHY, so as to quantify the improvements at the PHY BER and the link layer throughput – our results clearly illustrate the improvements STBC yield.
- Second, we incorporate the PHY implementation into the `ns` simulator – extending the `ns` with an *accurate* 802.11a PHY model allows us to quantify the impact of STBC at higher layers as well.

At first sight, it appears to be obvious that a method (such as STBC) which improves PHY performance will eventually have a positive impact on the overall network performance. However, there are several reasons which make the STBC impact on 802.11a an issue that warrants thorough investigation. Among them are the following questions:

1. STBC is a PHY technique, and the 802.11 standard was designed in such a way that the PHY is independent from the MAC [4, pg. 149]. How do PHY improvements reflect on the network

performance?

2. There is a complex interaction between PHY, link layer, MAC, and TCP. For example:

(a) Retransmissions at the link layer may be used so as to conceal PHY errors from TCP. If the PHY errors are reduced because of STBC, will the TCP throughput significantly increase?

(b) In heavily-loaded scenarios, packet collisions are known to be the limiting performance factor – will STBC improve performance in such scenarios?

The main point that we attempt to raise in this paper is that though STBC is a PHY technique, its effects can be traced both at the link layer, and at the TCP. In order to raise this point, we employ a comprehensive, across-the-layers study: we start by looking at the PHY, then at the link-layer, and finally at the MAC and TCP (we will shortly see how STBC significantly improves the performance at all layers). Unfortunately, due to space limitations, in the ensuing discussion at some points we have to sacrifice depth to breadth of coverage. PHY, MAC, and TCP are important research areas per se, but, in order to describe our comprehensive, across-layers study we cannot go thoroughly over all the technical details. *However*, we do try to explicitly state our modeling assumptions so as to delineate the accuracy of our approach and facilitate potential replication of our results.

The rest of this paper is structured is follows: in Section II we provide a brief overview of the 802.11 MAC, the 802.11a PHY, and we discuss how the 802.11a PHY may be enhanced with STBC. In Section III we describe our integrated approach to network simulation: we provide some details about the PHY simulation, and the extensions to the `ns`. In Section IV we present our PHY simulation results, and present a theoretical analysis on the impact of STBC on the link layer throughput. In Section V we investigate how the PHY and link-layer improvements improve the TCP performance in unloaded, and mildly loaded network scenarios. Finally, Section VI concludes the paper.

II. Model Description

Let us start by briefly revisiting the 802.11 MAC, the 802.11a PHY, and how STBC can be incorporated into the 802.11a.

A. 802.11 MAC

Complete details about the 802.11 MAC standard can be found in the draft standard [4] or in tutorial articles, e.g., [5]. In this section we overview some of the main characteristics of the 802.11 MAC with an emphasis on the points relevant to our simulation studies. We note that "the architecture of the IEEE 802.11 MAC is intended to be PHY independent" [4, pg. 149] – this motivates us to study the impact of STBC (a PHY technique) on the network performance.

802.11 Wireless LANs cover both ad hoc networks and infrastructure networks: in the latter type of networks, the access point (AP) serves as the equivalent of a basestation in a cellular architecture. In both types of networks, the basic channel access method is a variant of CSMA/CA (collision detection cannot be used, because of RF limitations, i.e., it is not possible for a wireless node to transmit to the channel and *at the same time* listen to it for any other ongoing transmissions). Collision avoidance is carried out in two ways: i) by following a *back-off* procedure before the transmission of data, and ii) by keeping track of the expected duration of ongoing transmissions[1]. The contention-based channel access method, which is called DCF (distributed coordinated function), assures fairness in the channel access, but it cannot provide strict QoS guarantees for time- or throughput-critical applications. Given the extensive body of work on CSMA/CA (which goes back more than 30 years), it is well documented that network performance deteriorates in high-load scenarios (too many packet collisions occur), and in the event of adverse channel conditions (which result in high BERs and lost packets). Consequently, even in the case of error free transmissions, the network performance (expressed in, e.g., overall throughput) is upper-bounded by the effect of collisions. In our present study, we are particularly interested in how STBC not only improve the wireless link quality, but the overall network performance: for example, in the Simulations section we will see that STBC improves performance even in overloaded networks.

On the other hand, contention-free communication can optionally be supported by the Point Coordination Function (PCF); the so-called Point Coordinator (PC), which normally is the AP, essentially imposes a TDMA-like channel access mechanism, where channel bandwidth is allocated by scheduling or polling. Important for our simulations later on is the observation that

[1] The latter technique is termed *virtual carrier sensing* and essentially boils down to recording the duration information contained in request-to-send (RTS), clear-to-send (CTS), and data frames.

PCF can implement a dedicated link between two nodes: this dedicated link is reminiscent of a connection-oriented link in TDMA systems. Herein, we also focus on how STBC improves performance for two nodes which communicate over such a "dedicated link", and we will shortly see that STBC presents to TCP a "smother" wireless link, which results in significant network performance improvements. Finally, we note that both PCF and DCF can coexist by simply alternating between CFP and CP.

B. 802.11a PHY

In this paper, we focus on the 802.11a PHY, which is OFDM-based and operates at the 5GHz band. The advantages of OFDM transmissions over the wireless medium are well-documented (see, e.g., [6]): OFDM is resistant to multipath (frequency-selective) fading, and is amenable to efficient FFT-based equalization (the equalizer is just a scalar multiplier at each OFDM subcarrier in the frequency domain). In the Appendix, we provide a short description of an OFDM system.

TABLE I

TRANSMISSION MODES

Mbit/s	Modulation	Coding rate
6	BPSK	1/2
9	BPSK	3/4
12	QPSK	1/2
18	QPSK	3/4
24	16-QAM	1/2
36	16-QAM	3/4
48	64-QAM	2/3
54	64-QAM	3/4

Though details can be found in the standard, herein we highlight aspects of the PHY that are directly related to our study. 802.11a supports a "data payload communication capability" of 6, 9, 12, 18, 24, 36, 48, and 54Mbit/s. These transmission rates are obtained by varying the modulation type or the channel encoding rates: these are depicted on Table I (corresponding to

[7, Table 78]). For our study, we need to emphasize the term "data payload" and point out the fact that the actual data throughput is a function of both the transmission mode *and* the channel conditions. For example, the 6Mbit/s mode is more robust to adverse channel conditions (such as noise, interference, and severe multipath) than the 24 Mbit/s mode: consequently, the 6Mbit/s can be actually faster than the 24Mbit/s mode because of potentially fewer lost packets[2]. This relationship between transmission rate at the PHY and link-layer throughput is well-established – it is also clearly illustrated in our simulation results of Section IV, and it has a fundamental impact on how the system should be optimized for data throughput maximization.

Of interest to the ensuing discussion of how STBC improves performance are the measures that the 802.11a PHY takes so as to combat harsh fading channel conditions. As specified in [7, pg. 54], before the data symbols are transmitted, the following operations are applied to the binary data stream:

1. scrambling, which removes spectral lines from the data stream
2. encoding with the "industry standard" rate 1/2 convolutional code with generator polynomials $g_0 = 133_8$, $g_1 = 171_8$
3. puncturing (which depends on the transmission mode)
4. interleaving, which is done in two steps. In the first step, consecutive coded bits are mapped to nonadjacent subcarriers; the second step maps consecutive coded bits onto the less and more significant bits of the constellation [7, pg. 17].

Furthermore, the OFDM transmission becomes more robust with the use of a relatively long cyclic prefix[3], the insertion of pilot tones, the use of short/long training sequences, and the use of null subcarriers (located at both ends of the transmission spectrum).

Let us now lay down the foundation of our work, by looking at how STBC can be incorporated in the 802.11a PHY.

C. STBC-enhanced 802.11a PHY

As the 802.11 standard clearly dictates that control frames should be understood by all nodes [4, pg. 95], and, in order not to raise any backward compatibility issues with existing equipment,

[2] As it is well known, in a single-hop Wireless LAN environment, packets rarely get lost because of congestion; rather they get lost because of channel errors or collisions.

[3] In the 802.11a terminology, the cyclic prefix duration is denoted as T_{GI}, GI stands for "guard interval", and $T_{GI} = T_{FFT}/4$, where T_{FFT} is the duration of an OFDM symbol [7, pg. 9].

we will constrain our development in the transmission/reception of DATA packets only (though our results can be extended without undue complications to CONTROL/MANAGEMENT packets as well). Hence, we assume that once a connection has been established between mobile nodes (or between mobile nodes and the Access Point (AP)[4]), then STBC may be used to boost the network performance. Naturally, the communicating nodes need to have STBC capabilities: the sender node should have more than one transmit antenna, and the receiver should be able to decode STBC signals. Given power/size limitations of mobile nodes, the nodes with multiple transmit antennas would naturally be the APs. Hence, by equipping the AP with more than one transmit antenna, and using STBC to transmit DATA packets, the quality of the wireless link can be significantly improved. Furthermore, the small decoding complexity of STBC codes [2] makes it possible to enhance the mobile nodes with STBC-decoding capabilities as well.

Such a scheme improves mainly the downlink performance[5]. It is certainly envisioned that in many scenarios, downlink will be the bottleneck (e.g., the transmission of a digital movie in a home entertainment system). On the other hand, as hardware costs drop down, it could be the case that it will be economically feasible to manufacture mobile nodes with STBC transmitters.

In the Appendix we provide a short mathematical description of how the 802.11a PHY can be enhanced with STBC. Let us now discuss our approach to simulating 802.11a Wireless LANs while taking into explicit account the PHY.

III. INTEGRATED APPROACH TO NETWORK SIMULATION

A Wireless LAN is quite complicated a system, and an accurate, across-all-layers simulation appears to be a challenging task. In this section, we discuss our simulation approach which consists of two major components: i) the PHY simulation (which produced PHY traces, BER and link-layer throughput plots), and ii) the network layer simulation (which is based on the `ns` simulator, and the extensions we made to its wireless model).

A. 802.11a STBC-enhanced PHY simulations

We have implemented in MATLAB the 802.11a PHY standard in the discrete time complex baseband equivalent model. As briefly described in the Appendix, we have also enhanced the

[4]As we mentioned earlier, in 802.11, the AP could be thought of as the equivalent of a basestation in a cellular architecture.

[5]Note though that the uplink performance could be improved with array processing (e.g., with maximum ratio combining at each OFDM subcarrier).

transmitter/receiver with the Alamouti [8] code, which is the most common member of the STBC codes [2]. To assess the PHY performance, we have used the channels of [9], which model typical Wireless LAN environments. The BER and link-layer throughput plots of Section IV (produced by Monte-Carlo simulations) depict the universal improvement achieved with STBC for all transmission modes of 802.11a. We underline that the MATLAB PHY implementation does not only produce BER or link-layer throughput curves, but it also generates packet transmission/reconstruction traces, which are fed in the `ns` simulator. Let us give an illustrative example of what we mean by transmission/reconstruction traces.

TABLE II

FRAME RECONSTRUCTION

SNR (dB)	802.11a	STBC+802.11a
35	Joy, bright spark of divinity, Daughter of Elysium, Fire-insired we trea	Joy, bright spark of divinity, Daughter of Elysium, Fire-insired we trea
15	`Hkm:;u3U8BF{n]b?L&X`	Joy, bright spark of divinity, Daughter of Elysium, Fire-insired we trea

Though BER plots is a commonly-adopted evaluation tool of PHY performance, at the network layer we are ultimately interested in whether a transmitted packet is received correctly. To assess the performance at the network/transport layer, we simulated the actual transmission of data frames (following the specifications in [7]). Reference [7, Annex G] presents an example of encoding a data message for the 802.11a PHY. The data message consists of the first 72 characters of the "Ode to Joy", and, quite naturally, the goal is to transmit, receive, and successfully reconstruct the transmitted frame. Table II depicts a representative snapshot of the reconstructed frame for two different channel qualities: at high SNR ("good" channel) the channel convolutional code is capable of correcting all bit errors that might have occurred, and the frame is successfully reconstructed. On the other hand, at relatively low SNR ("bad" channel), many more bit errors occur, the channel convolutional code is not capable of correcting all of them, and the frame cannot be reconstructed. We can see that in both cases, with STBC, the frame goes through without errors[6], whereas under 802.11a the frame was lost in the bad channel. Consequently, it is clear how STBC improves the wireless link quality, and results in significantly

[6] When more bit errors occur than what the FEC can handle, the CRC check at the receiver marks the frame as "lost".

fewer bit errors – a fact which is further corroborated by the plots of Section IV.

Using the aforementioned 802.11a PHY simulation model, we generated long traces of frame transmissions/reconstructions[7]. How these traces were incorporated into the `ns` simulator is the topic of the next section.

B. `ns`-Extensions

The wireless components of the `ns` [10] are primarily the result of the `monarch` project [11], [12] at the Carnegie Mellon University. According to the existing `ns` implementation in the absence of packet collisions, the packet error probability is a function of parameters such as power, node/antenna location, and channel statistics (e.g, Rayleigh or Rician fading). This channel modeling approach attempts to approximate the statistical behavior of the PHY at a low run-time computational cost. However, such an approach does not faithfully capture the characteristics of a specific transmission/reception technique at the PHY – this is why we had to extend the `ns` simulator in order to bring to the surface the impact of STBC.

Ideally, one would prefer to modify the `ns` simulator in such a way that whenever a packet was to be transmitted/received over the wireless link, the actual PHY event is simulated in `C++`[8]. Unfortunately, this is a rather time-consuming operation, which makes simulations rather long. Consequently, we opted to extend the `WirelessPhy` class (which is primarily responsible for the simulation of packet transmission/reception over the wireless channel), so that it can optionally read our MATLAB-PHY generated traces, and base packet-reception decisions upon them. Furthermore, we introduced an `ErrorModel`-descendant class, which simulates packet error events based on the aforementioned MATLAB-generated traces.

We emphasize that our trace-based approach allows us to simulate the actual wireless transmission/reception in a more faithful way, and enables us to draw conclusions on the impact of STBC on network/transport performance.

C. Assumptions on the simulation set-up

There is a number of assumptions we made in all simulation examples we present. For the PHY layer simulations, we assume perfect channel knowledge at the receiver, and perfect fre-

[7]Each trace is nothing but a file with 0's and 1's denoting successful frame reconstructions. Many such trace files generated corresponding to different SNRs and channel realizations.

[8]This entails proper modification of the `WirelessPhy::sendUp`, `WirelessPhy::sendDown` methods.

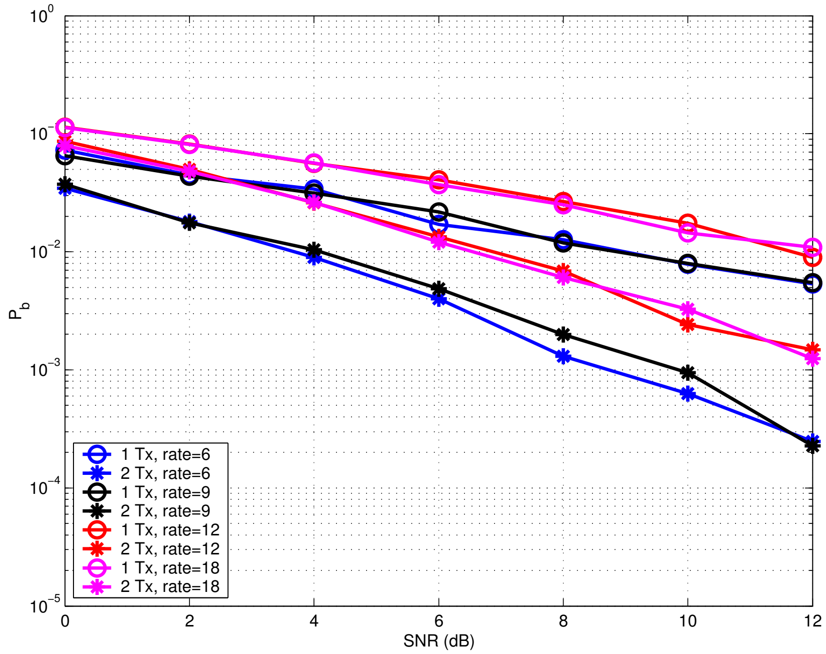

Fig. 1. BER for channel "A"

quency and timing synchronization. For the MAC simulations, we adopted assumptions similar to those that have appeared in, e.g., [5]: for example we assume that all nodes can hear one another (no hidden terminals).

IV. PHY AND LINK LAYER SIMULATIONS

Our "lower-level" 802.11a simulations consist of two components: the first component is related to the BER improvements due to the STBC-enhancements, while the second is component is related to the (ideal) link-layer improvements that STBC could yield.

A. BER improvements

In order the assess the performance of the 802.11a PHY and the improvements of STBC under a realistic set-up, we used two of the channel models of HiperLAN/2 [9]. The PHY specification of HiperLAN/2 is similar to that of 802.11a, and [9] contains information about channel models which are based on measurements in indoor/outdoor environments. Specifically, channel "A" corresponds to a typical office environment (with no light of sight) , which is modeled as an

FIR channel of order $L = 8$: each of the 9 independent taps of the channel is a complex normal variable (but there is a power profile). Channel "B" corresponds to an office environment with larger delay spread ($L = 15$).

Fig. 1 depicts the BER for 802.11a and the STBC-802.11a for the channel "A": the BER curves were obtained by generating a number of independent channel realizations (modeled after the "A" specification), and averaging their respective bit errors. We can clearly see the improvement of STBC transmissions in typical Wireless LANs environments. This is not a very surprising result, given that the merits of STBC transmission have already been analyzed in flat-fading channels, and an OFDM system (with sufficiently long interleaving) attempts to transform a frequency-selective channel to a group of flat-fading channels (which are the subcarriers in the frequency domain). Nevertheless, the aforementioned BER results further corroborate the merits of STBC *even under* the 802.11a specifications.

However, as we set out to demonstrate in this work, the BER improvements at the PHY manifest themselves as a performance-boost at higher layers. Before we discuss the results that we obtained using the `ns`, let us present some theoretical background.

B. Modeling the ideal PHY throughput performance

We underline that the throughput results we present in Section V were obtained using the frame transmission/reception traces of our PHY simulations. In this section, we provide theoretical intuition to explain the results that we present later on. We look at the case where a mobile node has been assigned (by the AP) a "logical" channel (see also, e.g., [13] and references therein). The user transmits frames of length L_p, which should be recovered without errors at the receiver side. Reliable reception is ensured by a hybrid FEC/ARQ scheme. The FEC is implemented by a rate R channel encoder; as a result, $L_p R$ information symbols are contained in every transmitted packet. The stop-and-wait ARQ utilizes an instantaneous, error-free feedback channel. At the receiver side the equalizer produces hard-decisions on the transmitted symbols. These decisions are fed into the channel decoder which produces estimates of the source symbols: if the decoder cannot recover successfully the transmitted frame, then the frame is retransmitted with probability p_r. The number of required transmissions until the frame is successfully received is geometrically distributed with average $(1 - p_r) + 2 \cdot p_r(1 - p_r) + 3 \cdot p_r^2(1 - p_r) + \ldots$

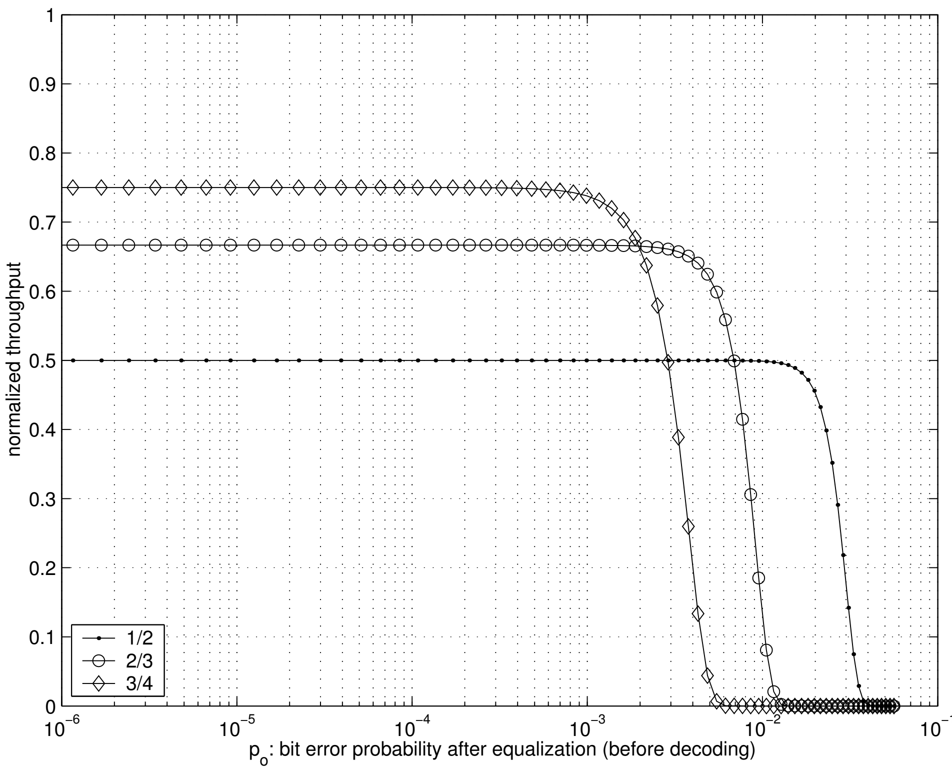

Fig. 2. Theoretical normalized throughput

$= 1/(1 - p_r)$ frame transmission time units. Consequently, the average throughput η is

$$\eta = L_p R(1 - p_r) \ . \tag{1}$$

Under the assumption that all errors are detected, p_r is given in terms of the information bit error probability p_b: $p_r = 1 - (1 - p_b)^{L_p}$. In our model, the decoder operates on hard-decisions which are provided by the equalizer. If p_o denotes the error probability of these hard-decisions (i.e, p_o denotes the transmitted bit error probability), then the information bit error probability p_b can be evaluated using well known techniques for the performance of channel codes in a binary symmetric channel with cross-over probability p_o (further details can be found in the Appendix). As a result, given BER curves (such as Fig. 1) a rough estimate of the link layer improvements could be obtained.

Fig. 2 depicts the theoretical estimate of the throughput as a function of p_o for various values of channel code rate R ($L_p = 1000$). It is evident that both the code rate (R) and the physical layer BER (p_o) have a tremendous effect on the expected throughput at the network layer. When

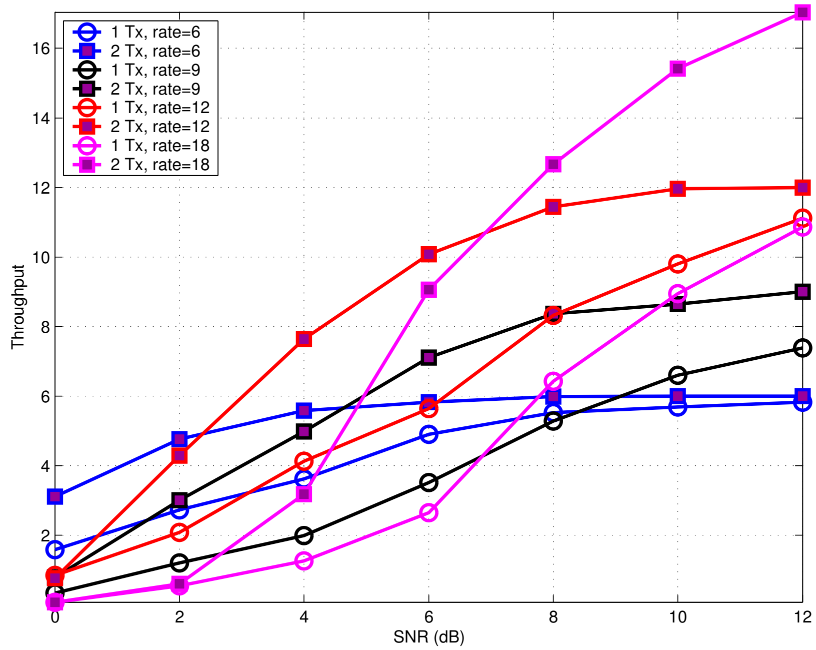

Fig. 3. Throughput for HiperLan/2 channel "A"

the physical layer cannot guarantee BER lower than 10^{-1}, then the rate-1/2 convolutional code offers the highest data throughput. However, once the BER is below 10^{-1}, then the rate-1/2 convolutional code hits the ceiling of 1/2 normalized data throughput. On the other hand, higher rate codes can yield higher throughput once the BER at the physical layer is below 10^{-1}. Interestingly enough, for higher rate codes, even the slightest improvement at physical layer BER is translated to considerable throughput gains at the network layer (note that figures similar to Fig. 2 have been reported in the past (see, e.g., [14] and references therein) and form the basis of adaptive ARQ systems which attempt to match the code rate to the underlying channel conditions).

These theoretical results are in agreement with throughput curves that we obtained by simulations. Fig. 3 depicts the link-layer throughput under channel "A" conditions There are some interesting observations that can be made from these figures

• The link-layer throughput gains with STBC are "universal". In almost all cases, STBC-enhanced 802.11a achieves a particular throughput value at a much lower SNR value than the 802.11a (in other words, STBC-enhanced 802.11a can afford worse channel conditions, and still

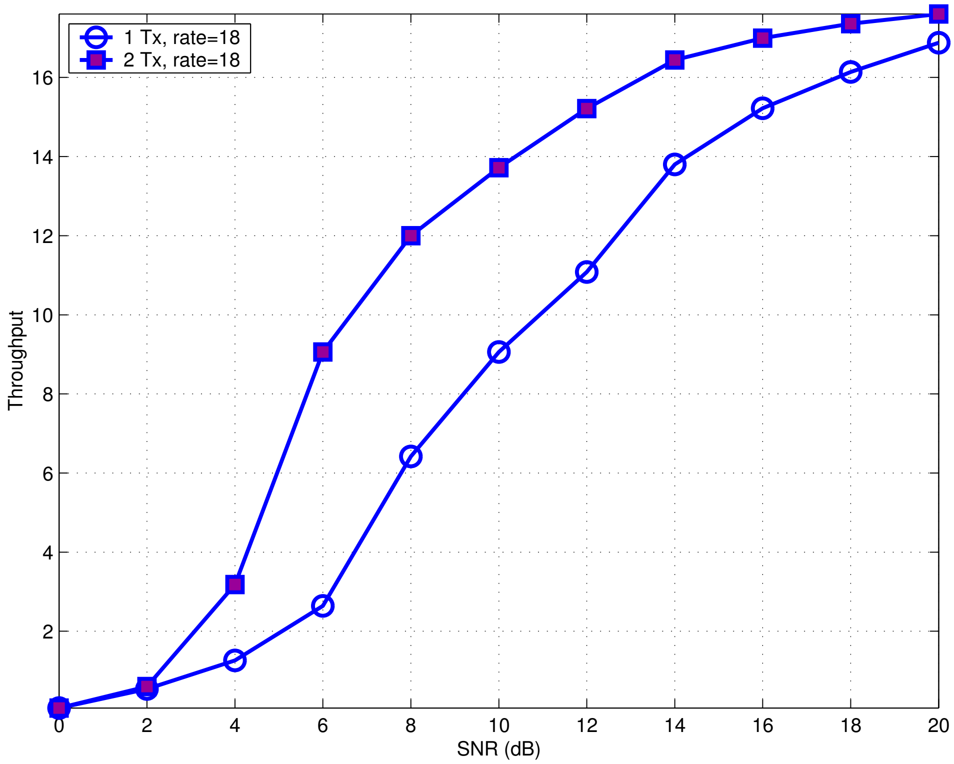

Fig. 4. Throughput for HiperLan/2 channel "A", 18Mbit/s transmission mode

result in the same throughput performance).

• *STBC modify the SNR region under which a particular transmission rate should be chosen.* For example, Fig. 3 indicates that without STBC it appears that one should transmit at 6Mbit/s when the SNR is 8dB, whereas with STBC it is possible to transmit almost *twice as fast* at 12Mbit/s. From this point of view, the impact of STBC on the network performance is rather *self-evident*.

• *STBC increase the transmission range and improve robustness of Wireless LANs*. Fig. 3 points out that, with space-time coding, a particular throughput value can be achieved over a wider range of SNR. Consequently, not only the transmission range can be increased, but also the robustness of the network (in the events of sudden fading) can be improved. Quantification of the STBC impact on robustness and reliability of 802.11 Networks is an issue that warrants thorough investigation in the future.

Though it is clear that STBC improves link-layer throughput, the question that naturally arises is how these results can be used to predict the performance at higher layers. We will see in the next section that when the link-layer allows retransmission of frames (so as to conceal the event of lost packets from TCP), the TCP performance plots resemble the throughput curves of this section. This perhaps surprising result, is due to the fact that TCP attempts to utilize the available "capacity" of a link: if STBC improve the performance at the link-layer, and catastrophic events (for TCP) are significantly reduced via link-layer retransmissions, then STBC should improve TCP performance as well (at least in the single-user case). This phenomenon is one of the topics explored in the next section.

V. MAC AND NETWORK LAYER SIMULATIONS

The theoretical and simulation-based results of the previous section allude to a potential performance improvement at the network layer. In this section we investigate the performance improvement for both the single and the multiple node/session case. Before we present simulation results, let us briefly discuss the cases where STBC is expected to have an impact on the network performance.

As we revisited in the previous sections, in Wireless LANs the performance drops in the presence of harsh fading channel conditions (which cause frame errors and packet drops). In the single-session case, if the link layer does not attempt to "hide" the frame errors from TCP then, in the presence of frame errors, TCP deteriorates fast, and effectively breaks down when the frame error rate is high enough. In this particular case, STBC is expected to yield an impressive performance improvement: STBC presents to TCP a much "smoother" channel. In Section V-A we see that in extremely harsh channel conditions, STBC result in *orders of magnitude* network performance improvement.

When the link layer does attempt to hide frame errors from the higher layers (and such is the case in 802.11), assessing the performance improvement induced by STBC becomes a more complicated matter. Link layer techniques do improve TCP performance and effectively decrease the probability that TCP will break down. In such a set-up STBC is expected to improve performance, because i) the number of frame retransmissions at the link layer is reduced, and ii) TCP will attempt to transmit at a rate equal to the perceived "capacity" of the wireless link.

The situation becomes a little bit more complicated in the presence of many nodes/stations contending for channel access under DCF. Let us assume that the system is moderately loaded, and the underlying physical channels for each of the communicating pairs are independent of each other. Then, even if a node experiences a particularly bad channel (and, as a result, the node

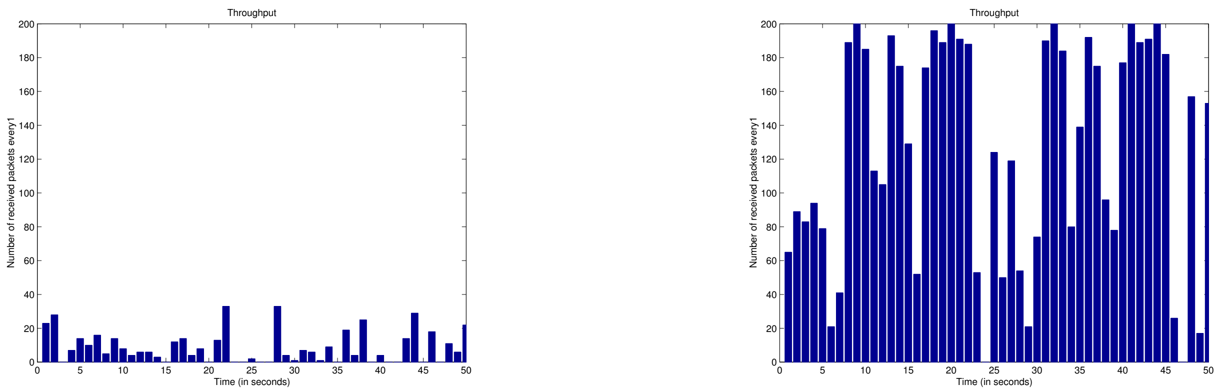

Fig. 5. Order of magnitude throughput improvement

backs-off its transmission until the channel is perceived to be good again), there will probably be another node ready to "jump in", and start transmitting. Consequently, one could expect that STBC do not yield a significant performance increase (measured as the aggregate throughput in the system). However, we will see that even in moderately loaded networks, the aggregate throughput is improved with STBC. Finally, even in over-loaded scenarios, where packet collisions at the MAC are the limiting performance factor, STBC still yield network performance improvements.

A. *Cases where STBC result in impressive improvements*

It is well-known that errors at the wireless physical layer reverberate across layers and have a negative impact on TCP performance (see, [15], and references therein). Roughly speaking, TCP interprets frame/packet losses as signs of network congestion, and cuts the transmission rate to a half, whenever these error events occur. When the link layer does not "hide" frame errors from TCP, TCP essentially breaks down. In fact, it has been shown in [16], that "TCP throughput deteriorates significantly when. . . TCP sees a probability of end-to-end loss above a threshold that scales as the inverse square of the bandwidth-delay product."

Fig. 5 illustrates a case where with only one antenna, TCP has essentially broken down, but with space-time block-codes, TCP throughput is significantly higher. For this particular example, we have assumed that there are no frame retransmissions at the link-layer, and that we operate at an SNR region where the BER performance (as given by Fig. 1) with only one transmit antenna is below the TCP breaking-point threshold, but the BER with STBC is above this particular threshold. Furthermore, it can be seen from Fig. 5, that with only one transmit antenna,

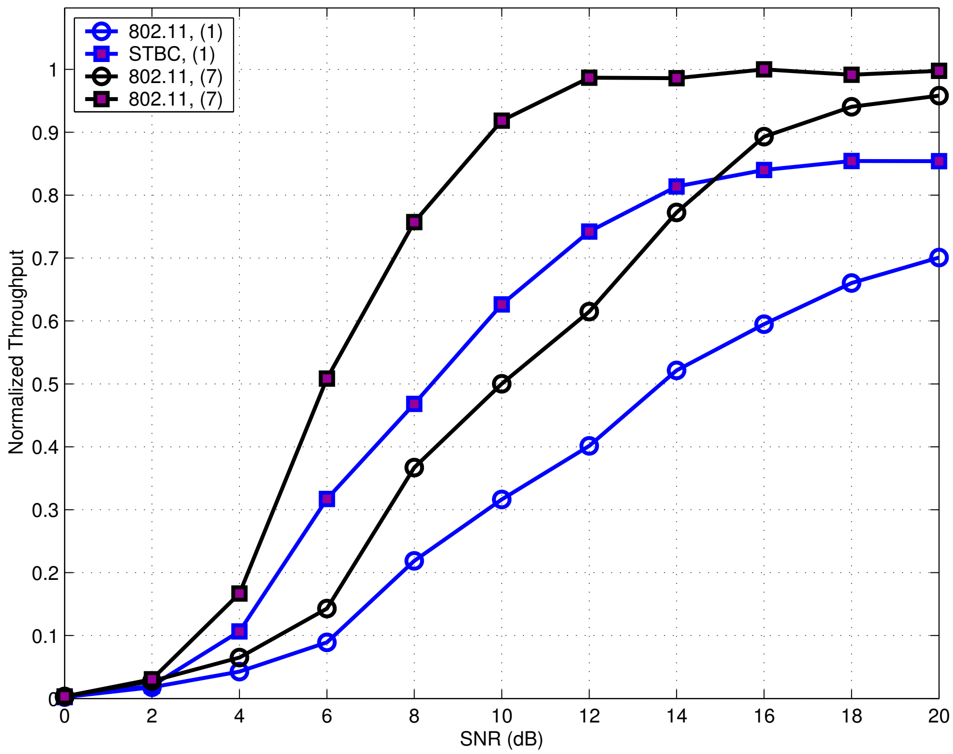

Fig. 6. Single Session

TCP times out, and does not transmit anything for a significant amount of time – however, this is not the case with STBC transmissions. Hence, the improvements with STBC are self-evident. It is also important to point out that STBC do not only shift the TCP breaking-point SNR threshold, but also STBC modify the statistics of the underlying frame-loss events[9] (see, e.g., [18] and references therein about how the statistics of frame losses affect TCP performance). These result illustrate that STBC essentially present to TCP a "smoother" channel, thereby improving the overall network performance.

B. DCF Link-layer: single TCP connection

Let us now introduce the "full" 802.11 protocol stack with DCF as the channel access method, and a link-layer with frame retransmissions. We assume a Wireless LAN with only two nodes active. In the first case, the transmit node has only one antenna and uses the 802.11a PHY, whereas in the second case the transmitter has two antennas and uses the STBC-enhanced 802.11a. Fig.

[9]For example, [17] discusses how STBC modify a Rayleigh fading channel to one with a chi-square distribution, which has shorter tails.

6 shows the achievable average throughput as a function of SNR. As it is expected, when the channel quality is sufficiently good, the difference between STBC-802.11a and 802.11a becomes smaller (however, with STBC-802.11a we would be able to switch to a faster transmission mode, as the throughput curves of Section IV-B indicate – this capability was not modeled in the simulation example we just presented). Also, it can be seen that as the number of allowable retransmissions at the link-layer is increased from 1 to 7, the throughput performance also increases. Nevertheless, the throughput with STBC is still significantly higher than the throughput with only one transmit antenna.

C. DCF Link-layer: multiple simultaneous TCP connections

Let us look at the multiple node case. We assume that many TCP connections (which transfer `ftp` data – the `ftp` sources are continuously transmitting) contend for the channel using DCF. In order to compare the 802.11a, and the STBC-enhanced 802.11a we look at two "extreme" cases: in the first case all nodes use the 802.11a PHY, whereas in the second all nodes have STBC capabilities. Fig. 7 depicts the aggregate throughput under the two schemes. Similar to the previous example, STBC yields network performance improvement throughout the SNR region, but especially at low SNR values. One point which is worth emphasizing is that both Fig. 6, and Fig. 7 resemble the link-layer throughput figures of Section IV-B.

Finally, we take a look at a heavily-loaded network scenario: 40 TCP connections which contend for the wireless channel. However, even in this case where many collisions occur, STBC result in improvement of throughput as Fig. 8 illustrates.

VI. CONCLUSIONS

In this paper we have presented an across-layers study of the STBC impact on the network performance in 802.11a Wireless LANs. Building upon an accurate simulation of the 802.11a PHY, we extended the `ns` simulator so as to investigate how a PHY technique (such as STBC) affects the network throughput. Through an extensive set of simulations, we have found that STBC improve performance at the PHY, link layer and TCP. Essentially STBC presents a "smoother" channel to TCP, a phenomenon which was also corroborated by a brief link-layer theoretical analysis. The results of our work serve as further testament to the potential of STBC as a building block towards faster, more reliable Wireless LANs. Additionally, the tools we developed

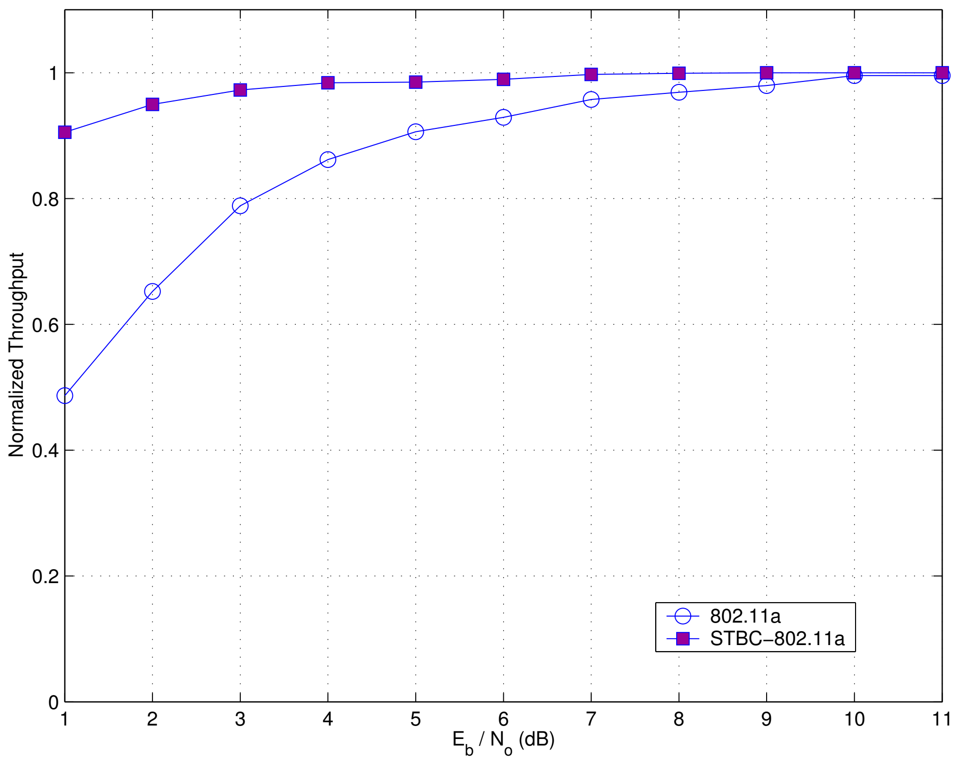

Fig. 7. 5 TCP sessions (300-byte frame)

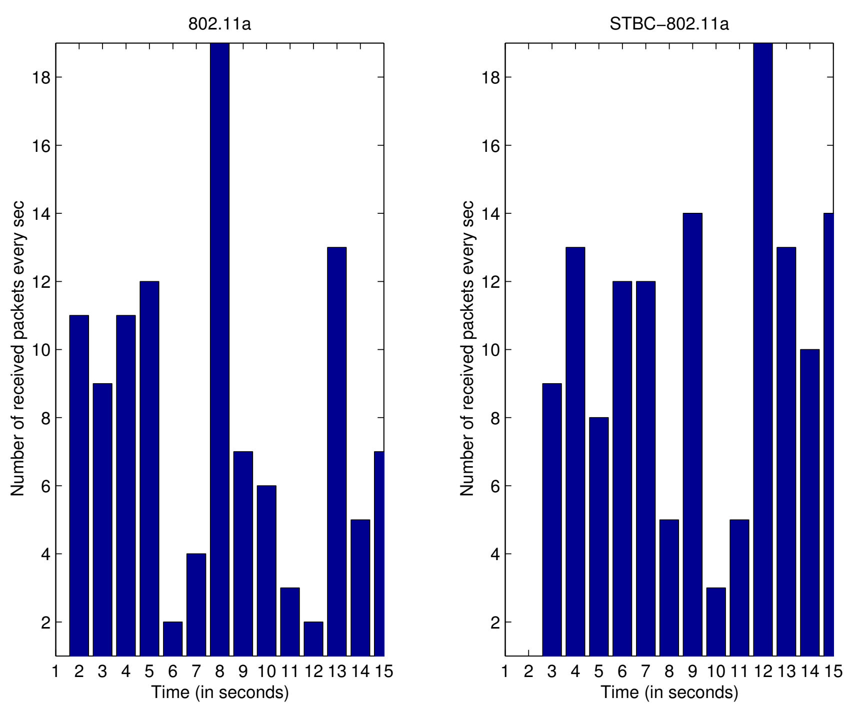

Fig. 8. Throughput of one (out of the 40) TCP sessions (at 7dB)

(STBC 802.11a PHY simulations, and `ns` extensions) could be used for a more accurate analysis and simulation of Wireless LANs.

APPENDICES

I. STBC AND STBC-ENHANCED 802.11A PHY

Herein we present a short description of OFDM, STBC, and an STBC-enhanced 802.11A PHY.

OFDM In OFDM transmissions, the information symbols X are grouped in blocks $\mathbf{X}$ of size N (see, e.g., [19]), and at the transmitter the IDFT is applied to them: $\mathbf{x} = \mathbf{Q}^{\mathcal{H}}\mathbf{X}$, (where $\mathbf{Q}$ denotes the $N \times N$ FFT matrix). Before the transmission over an order-L FIR channel $h(n)$, a cyclic prefix of length at least L is inserted into the data, which is discarded at the receiver. The received block $\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n}$ goes through a DFT operation to produce the frequency domain decision variables $\mathbf{Y}{:=}\mathbf{Q}\mathbf{y} = \mathbf{D}_H\mathbf{X} + \mathbf{N}$, where $\mathbf{D}_H$ is a diagonal matrix holding in its main diagonal the frequency response of the channel at each of the N subcarriers.

STBC With two transmit antennas, Alamouti's Space-Time Block Code [8] groups the input symbols into groups of two symbols, $\mathbf{c} = (c_1 \;\; c_2)^T$, which are fed to the space-time block encoder:

$$\begin{pmatrix} c_1 \\ c_2 \end{pmatrix} \longrightarrow \begin{pmatrix} c_1 & c_2 \\ -c_2^* & c_1^* \end{pmatrix} \begin{array}{l} \rightarrow \;\; \text{time} \\ \downarrow \;\; \text{space} \end{array} \tag{2}$$

With h_1, h_2 denoting the two coefficients corresponding to the flat fading channels from each of the two transmit antennas to the receive antenna, the received signals (for two consecutive time slots) are: $r_1 = h_1c_1 + h_2c_2 + n_1$, $r_2 = -h_1c_2^* + h_2c_1^* + n_2$, which can be written as:

$$\mathbf{r}{:=} \begin{pmatrix} r_1 \\ r_2^* \end{pmatrix} = \mathbf{H}\mathbf{c} + \mathbf{n}\;,\; \mathbf{H}{:=} \begin{pmatrix} h_1 & h_2 \\ -h_2{}^* & h_1{}^* \end{pmatrix}\;, \tag{3}$$

with n_1, n_2 independent AWGN processes each with variance $N_o/2$ per dimension. Under the assumption of constant channels (over the transmission of two symbols), it is easily seen that $\mathbf{H}$ is orthogonal, i.e., $\mathbf{H}^*\mathbf{H} = (|h_1|^2 + |h_2|^2)\mathbf{I}_2$. With $\tilde{\mathbf{r}}{:=}\mathbf{H}^*r$, $\tilde{\mathbf{n}}{:=}\mathbf{H}^*\mathbf{n}$, the transmitted symbols can be obtained from: $\tilde{\mathbf{r}} = (|h_1|^2 + |h_2|^2)\mathbf{c} + \tilde{\mathbf{n}}$.

STBC-enhanced 802.11a PHY The fundamental mechanism is to combine OFDM with STBC[10]. Two consecutive blocks $\mathbf{X}_1$, $\mathbf{X}_2$ yield the following $2N \times 2$ code matrix ($*$ denotes

[10] Combining OFDM with STBC was first proposed in [20].

conjugation):

$$\begin{pmatrix} \mathbf{X}_1 \\ \mathbf{X}_2 \end{pmatrix} \longrightarrow \begin{pmatrix} \mathbf{X}_1 & \mathbf{X}_2 \\ -\mathbf{X}_2^* & \mathbf{X}_1^* \end{pmatrix} \begin{array}{l} \rightarrow \quad \text{time} \\ \downarrow \quad \text{space} \end{array} \tag{4}$$

Each column corresponds to an OFDM block duration ("time"), whereas each row holds the symbols to be transmitted by each antenna ("space"): for example, during the second time OFDM block duration, $-\mathbf{X}_2^*$ is fed to the IFFT, a cyclic prefix is added, and it is transmitted by the first antenna.

At the receiver, after discarding the cyclic prefix, and taking the FFT, the received symbols in the frequency domain over the two consecutive OFDM block durations are: $\mathbf{r}_1 = \mathbf{D}_{H_1}\mathbf{X}_1 + \mathbf{D}_{H_2}\mathbf{X}_2 + n_1$, $\mathbf{r}_2 = -\mathbf{D}_{H_1}\mathbf{X}_2^* + \mathbf{D}_{H_2}\mathbf{X}_1^* + n_2$, which can be written as:

$$\mathbf{r} := \begin{pmatrix} \mathbf{r}_1 \\ \mathbf{r}_2^* \end{pmatrix} = \mathbf{H} \begin{pmatrix} \mathbf{X}_1 \\ \mathbf{X}_2 \end{pmatrix} + \mathbf{n} \, ,$$

which leads to space-time diversity gains at each subcarrier.

II. Theoretical performance

The information bit error probability p_b of a convolutional code of rate $R = k/n$ is upper bounded by (see, e.g., [21, pg. 490-491]): $p_b < (1/k) \sum_{d=d_{\text{free}}}^{\infty} \beta_d P_2(d)$. The coefficients β_d constitute the weight spectrum of the code, whereas $P_2(d)$ is the probability of selecting the incorrect path with distance d from the all-zero path. Using the Viterbi decoding algorithm on a binary symmetric channel (which can be used to model hard equalization decisions), an upper bound on $P_2(d)$ is ([21, eq. (8.2-31)]): $P_2(d) < (4p_o(1-p_o))^{d/2}$, with p_o denoting the crossover probability of the binary symmetric channel. The mother 1/2 convolutional code has 64-states, and generator polynomials $g_1 = 133_8$, $g_2 = 171_8$. Using the weight spectra (which can be found in [22, pg. 1117]), we can calculate an upper bound of the information bit error probability as a function of p_o (which is the transmitted bit error probability, i.e., the probability of bit error after equalization, but before the channel decoder).

References

[1] S. Lu, V. Bharghavan, and R. Srikant, "Fair scheduling in wireless packet networks," *IEEE Transactions on Networking*, vol. 7, no. 4, pp. 473–489, August 1999.

[2] V. Tarokh, H. Jafarkhani, and A.R. Calderbank, "Space-time block codes from orthogonal designs," *IEEE Trans. on Information Theory*, vol. 45, no. 5, pp. 1456–1467, July 1999.

[3] A.F. Naguib, N. Seshadri, and A.R. Calderbank, "Increasing Data Rate over Wireless Channels: Space-Time Coding and Signal Processing for High Data Rate Wireless Communications," *IEEE Signal Processing Magazine*, vol. 17, no. 3, pp. 76–92, May 2000.

[4] IEEE Std 802.11, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifations," ISO/IEC 8802-11: 1999(E), 1999.

[5] B.P. Crow, I. Widjaja, J.G. Kim, and P.T. Sakai, "IEEE 802.11 wireless local area networks," *IEEE Communications Magazine*, pp. 116–126, September 1997.

[6] Z. Wang and G.B. Giannakis, "Wireless Multicarrier Communications: Where Fourier Meets Shannon," *IEEE Signal Processing Magazine*, vol. 17, no. 3, pp. 29–48, May 2000.

[7] IEEE Std 802.11a, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifations. Amendement 1: High-speed Physical Layer in the 5 GHz band," Amd 1:2000(E), LAN MAN Standards Committee of the IEEE Computer Society, 2000.

[8] S.M. Alamouti, "A simple transmit diversity technique for wireless communications," *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 8, pp. 1451–1458, October 1998.

[9] ETSI Normalization Commitee, "Channel Models for HIPERLAN/2 in different indoor scenarios," Norme ETSI, document 3ERI085B, European Telecommunications Standards Institute, Sophia-Antipolis, Valbonne, France, 1998.

[10] L. Breslau, D. Estrin, K. Fall, S. Floyd, J. Heidemann, A. Helmy, P Huang, S. McCanne, K. Varadhan, Y. Xu, and H. Yu, "Advances in network simulation," *Computer*, vol. 33, no. 5, pp. 59–67, May 2000.

[11] A. Bonsall, D. Wang, and D. Stancil, "Propagation model embedded in a wireless network simulator," in *IMPAS 3rd Annual Wireless Communications Conference*, San Diego, CA, November 1998.

[12] R.J. Punnoose, P.V. Nikitin, and D.D. Stancil, "Efficient simulation of Ricean fading within a packet simulator," in *IEEE VTS Fall VTC 2000*, 2000, vol. 2, pp. 764 –767.

[13] A. Stamoulis, *Integrated Services in Wireless Networks with Jointly Designed Physical, Data, and Network Layers*, Ph.D. thesis, University of Minnesota, December 2000.

[14] R.H. Deng and M.L. Lin, "A type I hybrid ARQ system with adaptive code rates," *IEEE Trans. on Communications*, vol. 43, no. 2-4, pp. 733–7, Feb.-March 1995.

[15] H. Balakrishnan, VN. Padmanabhan, S. Seshan, and RH. Katz, "A comparison of mechanisms for improving TCP performance over wireless links," *IEEE/ACM Transactions on Networking*, vol. 5, no. 6, pp. 756–769, December 1997.

[16] T.V. Lakshman and U. Madhow, "The Performance of TCP/IP for Networks with High Bandwidth-Delay Products and Random Loss," *IEEE Transactions on Networking*, vol. 5, no. 3, pp. 605–615, June 1997.

[17] B. Hochwald, T.L. Marzetta, and C.B. Papadias, "A transmitter diversity scheme for wideband CDMA systems based on space-time spreading," *IEEE Journal on Selected Areas in Communications*, vol. 19, no. 1, pp. 48–60, January 2001.

[18] R. Zorzi, A. Chockalingam, and R.R. Rao, "Throughput analysis of TCP on channels with memory," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 7, pp. 1289–1300, July 2000.

[19] J.M. Cioffi, "A multicarrier primer," in *ANSI T1E1.4 Committee Contribution, No. 91-157*, Boca Raton,FL, November 1991.

[20] Z. Liu, G.B. Giannakis, A. Scaglione, and S. Barbarossa, "Decoding and equalization of unknown multipath channels based on block precoding and transmit-antenna diversity," in *Asilomar Conf on Signals, Systems, and Computers*, 1999, pp. 1557–1561.

[21] J. Proakis, *Digital Communications*, McGraw-Hill, 3rd Edition, 1995.

[22] D. Haccoun and G. Begin, "High-rate punctured convolutional codes for Viterbi and sequential decoding," *IEEE Trans. on Communications*, vol. 37, no. 11, November 1989.