

A Framework for Distributed Spatio-Temporal Communications in Mobile Ad hoc Networks

Gentian Jakllari ⁺, Srikanth V. Krishnamurthy ⁺, Michalis Faloutsos ⁺
Prashant V. Krishnamurthy [†] and Ozgur Ercetin ^{*}

⁺ University of California, Riverside, [†] University of Pittsburgh, ^{*} Sabanci University, Turkey
{jakllari, krish, michalis}@cs.ucr.edu, prashant@tele.pitt.edu, oercetin@sabanciuniv.edu

Abstract—Space-time communications can help combat fading and hence can significantly increase the capacity of ad hoc networks. Cooperative diversity or virtual antenna arrays facilitate spatio-temporal communications without actually requiring the deployment of physical antenna arrays. Virtual MISO entails the simultaneous transmission of appropriately encoded information by multiple nodes to effectively emulate a transmission on an antenna array. We present a novel *multi-layer approach* for exploiting virtual MISO links in ad hoc networks. The approach spans the physical, medium access control and routing layers and provides: (a) a significant improvement in the end-to-end performance in terms of throughput and delay and, (b) robustness to mobility and interference induced link failures. The key physical layer property that we exploit is an increased transmission range due to achieved the diversity gain. Except for space-time signal processing capabilities, our design does not require any additional hardware. We perform extensive simulations to quantify the benefits of our approach using virtual MISO links. As compared to using only SISO links, we achieve an increase of up to 150% in terms of the end-to-end throughput and a decrease of up to 75% in the incurred end-to-end delay. Our results also demonstrate a reduction in the route discovery attempts due to link failures by up to 60%, a direct consequence of the robustness that our approach provides to link failures.

I. INTRODUCTION

The use of antenna arrays in conjunction with space-time codes can significantly improve signal quality and thereby enhance the capacity of ad hoc networks. Depending on whether multiple transmitting antennas (inputs) and/or multiple receiving antennas (outputs) are used, one could have a Multi-Input Single-Output (MISO) system, a Single-Input Multi-Output (SIMO) system or a Multi-Input Multi-Output (MIMO) system¹. The deployment of antenna arrays on small mobile nodes, however, is infeasible due to the required size of these antennas. More specifically, the space between two elements of a multiple element antenna array must be at least of the order of $\frac{\lambda}{2}$, λ being the wavelength used for transmissions. For the commonly used 2.4 GHz frequency band, the required inter-element distance is 6.125 cm. Therefore, even an antenna with four elements can be too big to be mounted on a laptop and even more so on a PDA or a low cost sensor node.

A new paradigm that has emerged is the use of *virtual antenna arrays* (also called *cooperative diversity*) With cooperative diversity, nodes that are in the same vicinity simultaneously transmit *and/or* jointly receive appropriately encoded signals, i.e., the individual antennas on the multiple nodes are used together to

form an antenna array. With this method, one could create *virtual* MISO, SIMO or MIMO links. The use of virtual antenna arrays can yield the spatial diversity benefits possible with a traditional antenna array housed on a single node. While physical layer research on the use of virtual antenna arrays has been fairly extensive [36], [34], [35], [17], [30], there are no mature higher layer protocols, which can translate the advantages of using virtual antenna arrays to enhance network and application performance.

Our overarching objective in this paper is to translate the advantages of using virtual MISO at the physical layer into higher layer performance benefits. In this work, we define a virtual MISO link to be established when a group of nodes (transmitters) jointly enable space-time communications with a single receiver. For the virtual MISO link to be formed, the receiver needs to have an estimate of the channel state. We do not assume feedback i.e., the transmitters do not have any knowledge of channel state; the diversity benefits are achieved due to the use of space-time codes. The key advantage provided by a virtual MISO transmission is an *increase in the transmission range* with only a small increase in the channel interference. This is possible due to the improvement in signal quality on the virtual MISO link. However, exploiting this key benefit requires the establishment and use of such links in a networked setting. This is not trivial and requires significant changes at both the routing and the underlying MAC layers.

In this paper, we propose a multi-layer approach to exploit virtual MISO links in mobile ad hoc networks. Our approach is based on the development of a *synergy* between the layers of the protocol stack; lower layers export appropriate information and optimization “handles” to higher layers, while higher layers allow for the refinement of the performance parameters of lower layers. In particular, we take advantage of the extended range virtual MISO links to establish shorter paths, which in turn, leads to an increase in throughput and a reduction in latency. First, we develop a new MAC protocol that closely ties in with the underlying physical layer to enable virtual MISO links. In particular, the MAC layer facilitates coordination between the collaborating nodes that transmit jointly on a virtual MISO link. Second, we design a routing protocol that can construct a path with virtual MISO links. Our approach has two attractive properties: (a) it is completely decentralized and nodes do not need more than local (one-hop) information, and (b) it provides robustness to link failures due to both mobility and interference effects. The latter property is facilitated via a dynamic *anycast* mechanism (to be discussed later) for establishing virtual MISO links.

We perform extensive simulations with physical layer models that include fading effects to evaluate our approach. We observe

This work is supported in part by, the U. S. Army Research Office under the Multi-University Research Initiative (MURI) grant W911NF-04-1-0224 and the NSF CAREER Grant 0237920.

¹Traditional systems wherein nodes have a simple single antenna element are referred to as Single-Input Single-Output (SISO) systems.

that our schemes can successfully help form and exploit virtual MISO links. They provide a significant improvement in higher-layer performance in terms of the observed end-to-end throughput and delay. In particular, in mobile scenarios, the throughput increases by as much as 150% and the latency is reduced by up to 75% as compared with a traditional layered protocol stack which uses only SISO links. In addition, our approach results in a dramatic reduction in the number of route discovery attempts (by up to 60%) as compared with an existing on-demand routing protocol that is used over SISO links.

The rest of the paper is organized as follows. In Section II, we describe the relevant physical layer background; we highlight those features that influence our higher layer protocol design. We present our multi-layer approach with an emphasis on our design at the MAC and routing layers in Section III. The results from our simulations and a deliberation on the observations form Section IV. In Section V, we describe related work. Our conclusions and a discussion on possible future work are presented in Section VI.

II. PHYSICAL LAYER DEPENDENCIES

In this section, we provide a brief discussion on physical layer issues that are tied in to the protocols that we discuss in later sections. First, we present a brief overview of space-time codes and their impact on virtual MISO links. Second, we describe the impact of using multiple simultaneous transmissions on the transmission and interference range in ad hoc networks. Third, we discuss how channel estimates (critical for space-time communications) can be facilitated in practice via the use of pilot tones. Fourth we consider the impact of differences in average received power and delays of signals from multiple transmitters at the destination node. Finally, we briefly discuss other relevant issues.

Virtual MISO links and Space Time Codes: In a SISO system, the single transmitter would send m symbols in mT_s seconds for a symbol rate of $1/T_s$. On a virtual MISO link, there are N transmitters that transmit m complex symbols $\pm s_i, \pm s_i^*$ over kT_s seconds; here, s_i^* is simply the complex conjugate of the symbol s_i and $m \leq k$. In the presence of independently flat Rayleigh fading channels between the many transmitters and the receiver, this approach can provide large diversity gains if the symbols are transmitted in a particular pattern (called a space-time block code²). The receiver with knowledge of the complex channel fading coefficients h_i can linearly combine the multiple signals to recover the symbols with a much lower bit error rate (BER) than otherwise. The Alamouti code [2] is a well known example of space-time block codes with diversity of order 2 (two transmitters).

Diversity Gain: To get a sense of the diversity gain, consider a target BER of 10^{-3} . With the virtual MISO link under consideration, for this target BER, the required E_b/N_0 is 15 dB, whereas the needed E_b/N_0 is 25 dB on a SISO link [2] [11]; E_b refers to the energy in a bit and N_0 is the power spectral density of white noise. Thus, *with diversity gain, the signal can be recovered at a distance farther than the case where there is no diversity* (this

²Space-time block codes are characterized by a $k \times N$ matrix \mathbf{S} that specifies the *pattern* as per which symbols must be transmitted by the N antennas in each of the k time units of duration T_s . The rows correspond to *time* (the times at which the symbols are transmitted) and the columns to *space* (the antenna elements on which they are transmitted).

is a direct consequence of the SNR requirement for a given BER being reduced). The diversity gain is more pronounced for lower BERs. As the order of diversity (which is the number N of independently fading copies of the signal) increases, the diversity gain also increases.

Bandwidth and Power: In virtual MISO, the symbol rate will be $\frac{m}{k} \frac{1}{T_s}$. The measure of bandwidth utilization is the *rate* of the space-time block code $R = m/k$. If $m = k$, then, $R = 1$ (full-rate) and the bandwidth is completely utilized. It is more difficult to achieve full-rates ($R = 1$) with space-time block codes of higher orders of diversity. There are rate 1/2 and 3/4 codes that have been constructed in [40] that achieve higher orders of diversity ($N = 3$ and $N = 4$ transmitters). However, there is an associated penalty of lower bandwidth utilization (which we account for in our simulations). With the use of higher level modulation schemes, lower rate codes ($R < 1$) can still have a better performance than a full rate Alamouti code in terms of both bandwidth utilization and BER for a given SNR [11]. Each transmitter on a virtual MISO link need use only $1/N$ times the power P of a single transmitter on a SISO link [11]. Alternatively, it is possible to have each of the transmitters use the same power P ; in such a case, the power used on the virtual MISO link is N times the power used on the SISO link and thus, the range can be further enhanced.

Impact of Increased Transmission Power on Range and Interference: If the N transmitters forming a virtual MISO link, all transmit at a power P (the power used by a transmitter in the corresponding SISO case), the total transmitted power is now NP . Both the transmission and the interference range will increase. Nodes that are beyond the interference range of a traditional SISO transmission may *sense* the transmission. However, the increase in the sensing range due to this N -fold increase in power is not high as we demonstrate here. Consider N simultaneous transmissions, each at a power P_t . Let P_{sens} be the received power that is needed for a node to *sense* the transmission. P_{sens} and P_t are related by a path-loss that is typically of the form:

$$\begin{aligned} 10 \log_{10}(P_{sens}) &= 10 \log_{10}(P_t) \\ &- 10n \log_{10}(d) + C(f_c, G) \end{aligned} \quad (1)$$

where d is the distance between the transmitter(s) and the sensing node, n is the path-loss exponent and $C(f_c, G)$ is a constant that depends on the carrier frequency f_c and antenna gains G . Let the *new* carrier sensing range with the N transmitters that collaborate to form the virtual MISO link be d_1 . For simplicity, we assume that the transmitting nodes are in the close proximity of each other such that the distance from each of these nodes to the sensing node is approximately the same i.e., $\approx d_1$. The relationship now becomes:

$$\begin{aligned} 10 \log_{10}(P_{sens}) &= 10 \log_{10}(NP_t) \\ &- 10n \log_{10}(d_1) + C(f_c, G) \end{aligned} \quad (2)$$

From these two equations, we see that $d_1/d = 10^{\log_{10}(N)/n}$. Even with $N = 8$ transmitters, assuming $n = 4$ (commonly used as a baseline value in wireless channels [31]), $d_1 = 1.69d$. That is, the interference range increases by a factor of less than 2.

With a SISO link of transmitted power NP_t , the interference range would be identical to that in the above case (i.e., if $N = 8$,

the new range would be 1.69 times the range with $N = 1$). In addition, the transmission range would be extended only by the same constant factor; in other words, the transmission range of the SISO link would now be 1.69 times the transmission range with $N = 1$. However, with the virtual MISO link, the diversity gain enables us to increase the transmission range significantly compared to the increase in the sensing range. Let us suppose that the diversity gain achieved with the virtual MISO link (each transmitter using a power P_t/N) is D dB. The required SNR drops by D dB compared to a corresponding SISO link (with transmitted power P_t). If the range with the SISO link is d , the new range of the virtual MISO link will be $d_1 = d \times 10^{D/10n}$ where, n is the pathloss exponent. For $D = 10$ and $n = 4$, the range increase is by a factor of 1.77. Compared to a SISO link with transmit power P_t and range d , the virtual MISO range with each transmitter using P_t and diversity gain D would be $d_1 = d \times 10^{(10 \log_{10}(N)+D)/10n}$. For $D = 10$, $n = 4$ and $N = 8$, the range increases by a factor of 3.

Channel Estimation: For achieving the diversity gain, the receiver needs to have channel state information with respect to each of the transmitters. Note that the transmitters do not need to have this information (there is no feedback) simplifying the process of communication. The channel information can be derived by the receiver if the transmitters insert pilot symbols periodically (depending on how quickly the channel changes) [2] [7]. There will be some degradation if the channel estimates are not accurate. The pilot symbols from the N transmitters have to be orthogonal (they can be transmitted sequentially in time or made orthogonal in code) [2]. In the design of our protocols, we facilitate the transmission of these pilot symbols (or tones) so that the receiver can estimate the channel with respect to each of the transmitting nodes. The pilot symbols also help in recovering from synchronization errors (training the Viterbi decoder or the equalizer).

The pilot tones consist of a *known set* of symbols and we can assume that their *detection* is possible as long as the *average* power exceeds a certain threshold. In other words, even if these symbols were to experience harsh fades, they can be detected over fairly long distances. In our work, we assume that the pilot tones can be detected over the extended range of a virtual MISO link. Since the receiver expects to receive a known sequence, it can compute the channel coefficients, even if the sequence is corrupted due to fading.




Fig. 1. Scenario for relay and source node powers and relative delays

Relative Differences in the Power and Delay of the Received Signals: Since the cooperating transmitters are not co-located (as with a physical antenna array), the signals they trans-

mit are likely to be received at the destination node with different delays and average received powers (we assume that all nodes transmit at the same power level). Our discussion here is aimed at evaluating the impact of this effect on the performance of cooperative transmissions. For the purposes of this discussion, we use circular ranges, specific numerical (typical) values for distances between nodes, and average power values.

Consider Figure 1 where the source node is at a distance d from the destination. Cooperative relay nodes can be anywhere in a circle of radius R (typically = SISO range = 250 m) centered at the source node. If a relay node N_1 is at a distance $r_1 \leq R$ from the source node at an angle θ_1 with respect to the line joining the source and destination, its distance to the destination node is given by:

$$d_1 = \sqrt{d^2 + r_1^2 - 2dr_1 \cos \theta_1} \quad (3)$$

Similarly, a relay node N_2 at (r_2, θ_2) will be at a distance $d_2 = \sqrt{d^2 + r_2^2 - 2dr_2 \cos \theta_2}$ from the destination node. If the power of the signal $s_1(t)$ from N_1 at the destination node is P_1 , the power of $s_2(t)$ from N_2 will be (assuming $n = 4$):

$$P_2(\text{dB}) = P_1(\text{dB}) - 40 \log_{10}(d_1/d_2) \quad (4)$$

The relative delay between the two signals will be:

$$\tau = \frac{1}{c} \times (d_1 - d_2), \quad (5)$$

where, c is the speed of light. In the worst case, when the two relay nodes are diametrically opposite to one another and the destination node ($\theta = \pi$), the distance $d_1 = d - 250\text{m}$, $d_2 = d + 250\text{m}$, the relative delay will be $1.67 \mu\text{s}$ irrespective of d .

Now let us suppose that relay nodes are uniformly distributed in the circle shown in Figure 1. Our goal is to determine the cumulative distribution functions (CDFs) of the difference in received powers in dB and the relative delays between signals in seconds. These will give us an idea as to whether or not the average received powers will be drastically different and if there will be synchronization problems. Finding a closed form expression for the CDFs is difficult, but it is possible to compute the CDFs numerically for different values of d . The CDFs of the power difference and relative delays are shown in Figure 2. Only absolute values are shown in this figure (negative values of the power difference indicates that the signal from the second relay node is stronger than that from the first relay node).




Fig. 2. CDFs of (a) the power difference and (b) relative delays

An analysis of the CDFs reveals the following.

- 1) The relative delays between the signals are fairly small. In almost 80% of the cases, the delay difference is less than $0.6 \mu\text{s}$. If R is the raw data rate on the air and the modulation scheme accommodates k bits/symbol, the symbol rate will be R/k and the symbol duration T_s will be k/R . The symbol durations in 2 Mbps 802.11 WLAN (2bits/symbol, QPSK), 11 Mbps 802.11b (8 bits/symbol, CCK) and 802.11a/g (OFDM) are $1 \mu\text{s}$, $0.727 \mu\text{s}$ and $3.2 \mu\text{s}$ (with an 800 ns guard period). In all cases, there are physical layer approaches that can be used to combat the impact of lack of synchronization or *frequency selectivity* to this extent. We will discuss these approaches next.
- 2) The power difference can be substantial only if d is the same as the SISO range of 250m. Signals from one of the relay nodes can have a much higher power in this case because it can be much closer to the destination. We can however show that compared to the source node, the signal from any relay node has at most 12 dB lower power at the receiver even if $d = 250\text{m}$. As d increases (and these are really the scenarios where our protocols provide the most benefits), we can see that in more than 85-90% of the cases, the power difference between the signals from any two relay nodes will at most be 5 dB; in other words, the contributions from the two transmitters are significant in terms of achieving the overall diversity gain. *These results suggest that, in almost all cases wherein cooperative transmissions are used, the diversity gain is only dependent on the number of cooperating transmitters and not on the physical location of these transmitters (as long as they are all within the SISO range).*

Phase Synchronization between cooperating transmitters:

With cooperative diversity or Virtual MISO, it is often the case (as discussed above) that the propagation delays experienced by the signals from the different transmitters, en route the destination, are different. In addition, the clocks of the transmitters may not be perfectly synchronized. This leads to the *asynchronous* reception of the multiple signals. If two transmitters send their first and second symbols at the same times $t = 0$ and $t = T_s$, they may be received at $t = \tau_1, \tau_2$ and $t = T_s + \tau_1, T_s + \tau_2$. The symbol from the second transmitter received at time τ_2 might be interfered with, by the symbol from the first transmitter received at $t = T_s + \tau_1$. This effect is similar to what is seen with frequency selective channels i.e., inter-symbol interference (ISI) occurs. There have been several previously proposed physical layer techniques that can be used to overcome this problem [28], [5], [29], [24], [27], [16], [23], [41], [25]. We summarize these approaches briefly.

Time-Reversed STC: The first approach uses *time-reverse space-time codes* (TR-STC) [25], [22]. Here the symbols transmitted by one transmitter are reversed in time and transmitted by a second transmitter (this can be generalized for multiple transmitters as well). The use of TR-STC may result in a marginal reduction in data rate due to the requirement of a few *guard symbols* [27]. The Viterbi algorithm is applied at the receiver (with fairly low complexity) to perform a maximum likelihood sequence detection [22]. It has been shown in [27] that the BER performance even with large synchronization errors is very close to that under flat fading conditions and has minimal effect on data rates or receiver complexity.

Space-Time OFDM: A second approach is to use space-time

OFDM (ST-OFDM) where the frequency selective channel is converted into multiple flat fading channels each of narrower bandwidth. OFDM is already the modulation scheme of choice for 802.11a and 802.11g making ST-OFDM an attractive option. In both TR-STC and ST-OFDM, delays between signals that last for a few symbol durations can be handled with minimal penalty. In [27], the authors show, that, in cooperative communications ST-OFDM has a performance comparable to that with TR-STCs with an even simpler receiver implementation.

Use of an Equalizer: A third option is to treat the transmissions as being similar to *delay diversity* schemes [42] and use a decision feedback equalizer to achieve the diversity performance [41]. This approach may need some artificial delays to be introduced between the multiple transmitters.

Given these results, we assume in our simulations that the lack of synchronization and/or frequency selectivity are not problems and if they exist, straightforward techniques such as TR-STCs or ST-OFDM are employed to achieve the same diversity performance as synchronized space-time block codes under flat-fading conditions [27].

Impact of Doppler Spread: We further assume that the channel coefficients do not change over a few symbol durations (the channel fades slowly). At a carrier frequency of 2.4 GHz and a mobile station speed that is as high as 50 km/hour, the maximum Doppler spread is $f_m = 111.1$ Hz and the channel coherence time is $T_c = 1.61$ ms³. At a symbol rate of 1 Mega-symbols/s, the channel can be assumed to be constant for 16100 symbols. Thus, we make the reasonable assumption that the channel fading remains constant during the transmission of a packet.

Miscellaneous Issues: Note that in SISO wireless systems, error control coding and interleaving are employed to combat effects of Rayleigh fading. Thus, Rayleigh fading effects are commonly ignored in range calculations in ad hoc networks. Error control coding reduces the useful data rate and interleaving increases the decoding delay. A comparison between SISO links with such robust error control coding and virtual MISO links with space-time block coding is beyond the scope of this paper. In this paper, we have also ignored log-normal shadow fading that can impact the range in both SISO and MISO links. Shadow fading can be mitigated by increasing the transmit power [31]; we assume that the transmit power P_t is appropriately chosen to be high enough to overcome shadowing effects. We do not consider power control of any kind in this work. Power control could potentially reduce interference and differences in average received powers from cooperating relay nodes.

Finally, we wish to clarify that we do not propose any new physical layer or signal processing techniques. Our objective is to design and deploy higher layer protocols for efficiently utilizing the underlying physical layer capabilities.

III. OUR MULTI-LAYER APPROACH

We describe our cross-layer approach for establishing and exploiting virtual MISO links in mobile ad hoc networks. We highlight the interactions between the physical and the higher layers when needed.

Overview of the approach. Our approach can be decomposed into the following functional steps.

³The maximum Doppler spread is given by $f_m = f_c v/c$ where f_c is the carrier frequency, v is the mobile speed and c is the speed of light. The channel coherence time is approximately equal to $9/(16\pi f_m)$ [31].




Fig. 3. Routes that consist of virtual MISO links are depicted by solid arrows; links on the primary path are shown with dotted arrows. The A-H virtual MISO link is an alternative to the A-F link.




Fig. 4. Establishing a virtual MISO link.

- **Discovering the primary path:** Using existing routing protocols (such as DSR or AODV), the path constructed with SISO links between the desired source-destination pair is discovered. An example is shown in Figure 3; the path being referred to corresponds to the path A,B,C,F,G,H,I,J,L,M,N,O,P. We call this path the **primary path**, and it is: (a) the basis for establishing a path using virtual MISO links, and (b) the default solution in case virtual MISO links cannot be established.
- **Selecting relay nodes:** From among the nodes on the primary path, the nodes between which virtual MISO links are established, are identified. We call these nodes **relay nodes**. As we discuss later, we propose an intelligent *anycast* mechanism at the MAC layer to identify the relay nodes. A new route is now established via these relay nodes using virtual MISO links. In Figure 3, this new route would correspond to A,F,L,P (when each node is assumed to cooperate with 3 neighbors).
- **Forwarding data packets:** The relay nodes exchange data across the virtual MISO links and towards enabling this, we design a novel MAC protocol.
- **Increased robustness to link failures:** If any of the virtual MISO links were to fail, we resort to virtual MISO anycasts to reconstruct the route locally; only if this anycast were to also fail would the source generate a new route query request.
- **Improving the route dynamically:** Our approach attempts to improve the route (initially created) in a proactive and dynamic fashion. Given the opportunity, a node may *replace* its next hop relay with an alternate relay that can help create a shorter path. For example, in Figure 3, A might replace the virtual MISO link A-F with A-H, if it is possible.

Our design provides the following inherent advantages: (i) building our solution on top of SISO based solutions makes the system backward compatible without compromising on performance⁴ and, (ii) use of MAC layer anycasts provides an inherent robustness to mobility.

We elaborate on each of the functions listed above. In order to facilitate the discussion, we first present our MAC layer protocol and then discuss routing issues; note here that the two layers are tightly coupled.

A. Media Access Control using virtual MISO links

In order for a sender node to establish a virtual MISO link with a receiver node, it would need to elicit cooperation from its neighbors. In addition, the cooperating nodes would need to transmit pilot tones to enable the estimation of the channel state at the receiver, as was discussed in the previous section.




Fig. 5. The sequence of packet transmissions on a virtual MISO link.




Fig. 6. The effects of a virtual MISO link on non-participating nodes.

We develop a MAC protocol for creating and using virtual MISO links based on the principles of the IEEE 802.11 MAC protocol. The protocol can support virtual MISO anycasts which, as we see later, facilitates the selection of relay nodes. We describe the new MAC layer and its interdependencies with the physical and routing layers below.

Initiating the MISO transmission. Let us consider two relay nodes which wish to communicate via a virtual MISO link. The initiating node would *multicast* a modified *local RTS* message (transmitted by that node alone using SISO). The message contains the identities of the neighbor nodes⁵ which are invited to cooperate in a transmission. This process is illustrated in Figure 4. In the example, node S chooses the nodes A, B, C and H from among its neighbors. The node selects neighbors in a random

⁴We show later with simulations that our routing scheme constructs paths that are close to the best possible path in terms of hop count.

⁵We assume that each node is aware of its one hop neighborhood; HELLO messages are used to facilitate this. Note that this is a common assumption in previous works.

fashion as per a uniform distribution. This approach is reasonable since, as shown earlier, the achievable gain is dependent primarily on the number of cooperating transmitters. The choice is made on a per packet basis. This policy also has the advantage that it balances the *cooperation* load between the neighbor nodes across multiple packet transmissions.

Channel estimation and pilot tones. Upon the reception of the local RTS message, the chosen collaborating nodes would transmit pilot tones as was explained in Section II. It is critical that the pilot tones do not conflict with each other; they need to be orthogonal either in time or in code. In our implementation, we separate the pilot tones in time⁶. A simple rule, such as the node with the minimum ID transmits the pilot tone first and so on, can establish a transmission order. At the end of this process, if the receiver was available and able to detect the pilot tones, it will have channel estimates with respect to all of the collaborating nodes. It is possible that some of the chosen neighbors may be under the influence of interference (their NAVs indicate that they are busy) and thus, may not be able to cooperate. In such a case, the transmission fails (the source can hear the pilot tones and thus, can detect failure upon not hearing an expected tone) i.e., the source backs off. In other words, virtual MISO transmissions are allowed only if a fixed number of cooperating neighbors are available; this policy is followed to eliminate the possible effects of link asymmetry as will be discussed later.

Utilizing the virtual MISO link. Ensuing the transmission of the pilot tones, the cooperating neighbors would jointly transmit an **M-RTS message**⁷, which is the equivalent of an RTS message transmitted over a virtual MISO link. We depict the packet exchange process in Figure 5. Note here that bits that form the M-RTS packet are encoded using space-time codes as discussed in Section 2. Since the receiver, R in our example in Figure 5, already has channel state information (derived from the pilot tones), it can decode the M-RTS message. In response, node R needs to acknowledge the reception of the M-RTS, and for that, it needs to *collaboratively* send a message to S across a reverse virtual MISO link. We refer to the response message as the **M-CTS message**. Node R would first send a *local* CTS message to elicit the assistance of its own neighbors for a virtual MISO transmission. Following the local CTS, this new set of cooperating nodes (in the vicinity of R) will now send their pilot tones back to node S in the manner specified earlier for the M-RTS transmission. The pilot tones are then followed by the transmission of the M-CTS message. Note that the cooperative transmission of a CTS fail (as in the case of the RTS); if this were to happen, the communication would fail and a retransmission attempt will have to be made later.

After the reception of the M-CTS message, node S will instigate the transmission of the M-DATA, which is a data packet transmitted over the virtual MISO link (already established by the successful exchange of the M-RTS and M-CTS messages). First, the data packet is sent locally to the cooperating neighbors. Immediately after the local transmission, all the cooperating nodes and the S itself, will jointly send the M-DATA packet.

⁶Alternatively, one might separate the pilot tones by using spreading codes. Thus, the collaborating transmitters would transmit the tones simultaneously albeit with orthogonal spreading codes. The receiver could then use a correlation receiver to separate the transmissions as in a traditional CDMA system [31].

⁷For readability, we use the nomenclature from the IEEE 802.11 standard and use the prefix “M” to indicate that the transmission is over a virtual MISO link.

Finally, node R, with help from its neighbors (first a local ACK is transmitted), will perform the transmission of an M-ACK message across the reverse virtual MISO link. If the channel is slowly varying (as assumed above), the M-DATA and the M-ACK messages can be sent across the established virtual MISO link without additional pilot tones. If the channel is very dynamic, pilot tones may have to be inserted prior to M-DATA and M-ACK transmissions as well.

Effects of a virtual MISO transmission on other nodes. The creation and use of a virtual MISO link requires a careful consideration of the manner in which nodes that are not participating in the transmissions on a virtual MISO link, should update their network allocation vectors (NAV's). We modify the policies that are used with the traditional IEEE 802.11 MAC protocol to accommodate virtual MISO transmissions. To elucidate our approach, we first consider the example in Figure 6. In the figure, we show the transmission range of a SISO link, the sensing or the interference range due to a SISO transmission and the interference range induced by a virtual MISO transmission. Recall that these were discussed in Section II. Let us assume that the only transmission in progress is one that is initiated by the transmitter S and that R is the intended receiver. Let us assume that nodes D, C and F are chosen for the cooperative transmission on the virtual MISO link. Nodes A, B, and E (within the SISO transmission range) would update their NAVs upon the receipt of the local RTS message. Nodes G,H, I, J and K (within the SISO interference range) would update their NAVs after they decode the M-RTS message and thereby infer that they are not at the tail of the virtual MISO link being established. Nodes R, L, M, N, O, P, Q, T, U, V, W, X, Y and Z would detect the pilot tones and decode the M-RTS message. Except for R, the other nodes in this group simply go back to the “idle” state (NAVs are not updated), once they infer that they are not a part of the virtual MISO link (unless R seeks their cooperation explicitly). Later, if these nodes physically sense any of the virtual MISO transmissions (without channel estimates), they would update their NAVs.

Generalizing from the above example, our approach has the following effects on third-party nodes depending on their location.

Effects on Nodes within the SISO transmission range: Nodes that receive the transmission of a local RTS or a CTS message (over the SISO link) will set their Network Allocation Vectors (NAV) in accordance with the rules specified by the IEEE 802.11 MAC protocol. Specifically, if they are able to decode these messages and they *have not* been chosen by the source to cooperate on the virtual MISO transmission that is to follow, they update their NAVs to reflect that the channel is busy for the duration specified in the local RTS/CTS message.

Effects on nodes between the SISO transmission range and the SISO interference range: Nodes in this range, sense the carrier due to the local RTS/CTS messages (as they are in the SISO interference range). These nodes would anticipate forthcoming pilot tones and attempt to detect and decode these tones. The intended receiver *could* belong to this group of nodes. Nodes other than the intended receiver, upon decoding these virtual MISO based transmissions, will realize that the transmission is not for them (these messages carry the address and identity of the destination node). These nodes would then, update their NAVs to reflect that the channel would be busy for the next EIFS (Extended

Inter-frame Space) as specified by the IEEE 802.11 protocol.

Effects on nodes that are between the SISO interference range and the virtual MISO interference range: In this region, nodes do not sense the channel to be busy due to the transmission of a local RTS/CTS message. However, they are able to detect pilot tones and would subsequently attempt to decode the M-RTS or M-CTS packet. Again, note here that the intended receiver *might* be in this range. Except for the intended receiver (which would act as discussed previously), the other nodes in this region would simply leave their NAVs unchanged at this time. If later, they *physically sense* the M-DATA (or the M-ACK) transmission (now without channel estimates via pilot tones), they will update their NAVs to denote that the channel would be busy for the next EIFS.

Effects on nodes that are beyond the virtual MISO interference range: These nodes would recognize that neither are they required to participate in the transmission nor are they subject to any interference effects. These nodes can participate in other transmissions.

Handling packet losses and time outs. We handle packet losses using timers and a retransmission policy that is similar to that with the IEEE 802.11 MAC protocol. Every expected response has an appropriate time out, and a timer expiration is considered to be due to a packet failure; the node that times-out increases a failure counter. For example, if the unicast transmission of a packet on a virtual MISO link fails for a preset number of attempts, it is considered a link failure (of the virtual MISO link). In such a case, the MAC layer resorts to a local recovery using anycast transmissions as we discuss later. If this were also to fail for preset number of repeated attempts, then, a route discovery is invoked by the source. If virtual MISO links cannot be established, the primary path is used as the default path and SISO links are used.

B. Routing across Virtual MISO Links

Having explained the basic functions at the MAC layer, we now discuss the functions at the routing layer.

Identifying the primary route. As mentioned earlier, our approach can be built on top of any existing ad hoc routing protocol designed for SISO links. While we could have employed any traditional routing protocol for this purpose, we employ the popular dynamic source routing (DSR) protocol. A construction based on any other routing protocol (such as the Ad hoc On demand Distance Vector or the AODV protocol) would be similar. We first employ DSR to compute a primary path with SISO links. Once this primary path is discovered, we assume that the nodes on the path are aware of a chosen path identifier associated with the path. In addition, each node is aware of its relative position on the path or the distance (in number of SISO hops) from the source of the path from its routing table.

Selecting relay nodes via virtual MISO anycasting. The next goal of our routing layer is to “shorten” the primary path, or in other words, to substitute consecutive short-distance hops due to SISO links with long-haul virtual MISO links (Figure 3).

However, in selecting relay nodes, we impose the design constraint that the created virtual MISO links be bi-directional⁸. In

other words, if the sender on a virtual MISO link uses k collaborating nodes, the receiver must have at least k neighbors. As an example, in Figure 3, node H would be precluded from being a part of the virtual MISO route, since it does not have enough neighbors (in this example 3 neighbors) to reach to node A. Even though the basic requirement is that virtual MISO links be bi-directional, in our implementation, we use a more restrictive policy. We require that a node *must* collaborate with a preset fixed number of neighbors in order to participate in the formation/use of a virtual MISO link. This has the following advantages: (a) it ensures that the sensing range of all of the virtual MISO links in the network are the same (e.g. link asymmetries are avoided), and (b) it simplifies the MAC layer design by allowing for a *fixed* time allocation for the transmission of the sequential pilot tones as discussed earlier. The last advantage of our policy directly follows from the fact that a fixed number of nodes are expected to transmit pilot tones during any transmission over a virtual MISO link. A more dynamic and adaptive approach could be possible, but we do not explore it in this work.

The mechanism for selecting relay nodes should be both dynamic and distributed. With our approach, when the source or a relay node intends to perform the first packet transmission, it performs a **virtual MISO anycast** of the M-RTS message to the nodes on the primary path. The *anycast* could simply specify the label associated with the path that is made known to all the nodes during the route discovery phase.

When a node on the route successfully receives the M-RTS packet, it considers itself to be a candidate relay if it has a sufficient number of neighbors, as was discussed earlier. Naturally, if any of these nodes were busy (due to virtual or physical carrier sensing), they cannot respond to the M-RTS and thus, cannot be candidate relays. In order to avoid multiple responses to an M-RTS, the candidate nodes use a back-off timer, which is set in proportion to the candidate’s distance from the source. In our example in Figure 3, nodes B, C, and F hear the M-RTS transmission, and set time-outs as per their hop count distance from A, with F setting the shortest time-out. If the timer expires, the particular candidate node will respond with a *unicast* M-CTS transmission as discussed earlier. An M-CTS message sent as above would be overheard by the other competing candidates and these nodes would then abort their scheduled M-CTS transmissions. Once the next relay node is identified, it performs a similar anycast to identify the next relay until the final destination is reached.

After this initial construction of a virtual MISO path, for subsequent transmissions, we use **virtual MISO unicasts** of M-RTS messages between the identified relays; the M-RTS message is explicitly directed to the next relay. This policy alleviates the contention and back-offs that arise with anycasting. Note that only the M-RTS message is anycast during the first packet transmission; the M-CTS, M-DATA and the M-ACK are always unicast, i.e. are intended for a specific recipient node.

Aggressive versus conservative anycasting. We enable the source to control the aggressiveness of the anycast. Specifically, the source chooses a restrictive set of nodes on the primary path that can consider themselves as candidate relays. This restriction is imposed since having many candidate relays could result

⁸Without this requirement, the sender may be able to communicate with the receiver on the virtual MISO link; however, the receiver node may not be able to close the virtual MISO link and send acknowledgments to the sending node. One

way of coping with this would be to send such acknowledgments over multiple smaller virtual MISO hops from the receiver to the sender. This would introduce complexity, and could potentially degrade performance.

in collisions of either the pilot tones (on the reverse link) or the M-CTS messages. The restrictive set is chosen by imposing a limit on the distance (in hop count on the primary path) between consecutive relays on the virtual MISO path. The above restriction may lead to the construction of longer paths, since some nodes may be precluded from attempting a response even if they hear the M-RTS message. However, this is addressed with our *dynamic route reconfiguration mechanism* (which is discussed below).

Increased robustness to link failures. The anycasting approach allows for dynamic route recovery if any of the virtual MISO links were to fail. Links could fail if relay nodes were to move out of range or due to interference effects. After a protocol-specified number (as per the IEEE 802.11 specifications) of unicast retries, the transmitting node would regress to a virtual MISO based anycast to find an alternative relay node. Note that we do not recompute the primary path, but only perform an anycast to discover a new relay node. In other words, a virtual MISO based route could be reconstructed, even if some of the nodes on the primary path were to move out of range. As an example, in Figure 3, if the relay L were to move out of range, node F would perform an anycast and might possibly select node J as its next relay. Node J would then perform an anycast to successively reconstruct the route. The process drastically reduces the number of new route discovery broadcasts (necessary with traditional on-demand routing schemes). Thus, the use of anycasting over virtual MISO links also provides an inherent robustness to link failures.




Fig. 7. Improving the route dynamically: by overhearing H's transmission, A can choose to establish the A-H virtual MISO link.

Improving the route dynamically. We propose a dynamic route reconfiguration mechanism to improve upon an existing routing path. During data transfers, routes are continuously monitored and reconfigured to improve performance when possible. The key idea is to capitalize on the ability to overhear packet transmissions in order to identify opportunities for the construction of a better path. As an example, in Figure 7, the non-cooperative path is A,B,C,D,E,F,G,H,I,J, with A being the source and J being the destination. We assume that currently, the virtual MISO link A-D is enabled. Due to mobility⁹, node H might be later reachable via a virtual MISO link from A. A virtual MISO link from A to H would drastically reduce the end-to-end hop count. Thus, H is a better candidate than D for the establishment of a virtual MISO link from A. Our protocol is equipped with a reconfiguration mechanism to take advantage

⁹Alternatively, this could be the result of a conservative anycast policy as was discussed earlier.

of such opportunities. If node A overhears a transmission from H, it would recognize that H is reachable *and* it is further downstream from D. Node A, then, would attempt to establish a link directly to node H. The switch to node H is straightforward; node A begins to unicast packets to H instead of to D. An explicit notification of the change to node D is preferable; but even without such a notification, the routing information at D will eventually expire.

IV. PERFORMANCE EVALUATION

In this section, we describe our simulation experiments, report the results of the experiments and analyze the performance of our approach. We use OPNET [13] as our simulation platform.

A. Simulations Models and Parameters

Physical Layer Models: The current wireless channel model in OPNET accounts only for path-loss and noise. We modify the models to include flat Rayleigh fading. As discussed in Section II, we assume that the environment under consideration reflects the use of the 2.4 GHz band. Thus, we assume that the channel is slowly varying and does not change during a packet transmission.

We do not simulate symbol level transmissions as discussed in Section II. Instead, we choose to perform abstractions to reflect packet level effects. This drastically reduces the simulation time without significantly compromising on the quality of the performance results. Each SISO packet transmission is assumed to have an associated transmission range and a sensing range as with the IEEE 802.11 MAC protocol. Nodes that are within the transmission range could potentially receive the packet after a successful RTS/CTS exchange. Within this range the packet is subject to a random attenuation, chosen from a Rayleigh distribution to account for fading effects and if this attenuation drives the received SNR below a threshold¹⁰ (of SNR_{TH} dB), the transmission fails.

Computing the new range: In our experiments, all virtual MISO transmissions are invoked by five cooperating transmitters¹¹. We compute a new range (as described in Section II) derived from the corresponding diversity gain. This new range is computed as follows. If this diversity gain is D dB (as tabulated in [11]), the preset threshold SNR is now set to $(SNR_{TH} - D)$ dB. In our simulations SNR_{TH} and D are set to 25 and 15 dB, respectively. Then, with the path-loss model (a fourth order path loss exponent is assumed), we recompute a new transmission range based on this new lower threshold SNR as discussed in Section II. A given node could *potentially* communicate with any node within this new transmission range using a virtual MISO link.

Incorporating channel effects: During the construction of the virtual MISO link, the local RTS transmission is subject to the same effects as a SISO link. As discussed in Section II, we assume that the pilot tones can be *detected* (not decoded) over the

¹⁰The threshold would facilitate a target bit error rate.

¹¹As described earlier, our framework is applicable with any number of cooperating transmitters. For a different number of cooperating transmitters, the achievable diversity gain will differ. For the density of nodes chosen in most of our scenarios (average degree = 8), the number chosen (five) seems reasonable. Note that beyond a certain threshold, an increase in the number of cooperating nodes (antenna elements) is unlikely to yield a significant increase in diversity gain [11].

extended range of a virtual MISO link. The Rayleigh fade experienced by each pilot tone is recorded by potential receivers that do not sense the channel busy (if the pilot tones do not collide with other tones). Using the *channel knowledge* thus gained, each receiver is able to reconstruct the messages that are subsequently transmitted at the MAC layer. If α_i is the attenuation due to the Rayleigh fade suffered by the signal from the i^{th} transmitter, and d_i is the distance to the receiver from the i^{th} transmitter, the signal received by the receiver is now attenuated by a factor $\sum_{i=1}^5 \alpha_i^2 d_i^4$. We now compare the received SNR with the threshold ($SNR_{TH} - D$) and declare a successful reception if the former exceeds the latter. Additive white gaussian thermal noise is assumed¹². In our work, we assume that the data rates on the channel remain fixed at 2 Mbps¹³. The use of an appropriately chosen space-time code, given that there are five cooperating transmitters, leads to a utilization R of 0.8; in other words, 4 symbols are transmitted in 5 chosen time-units of duration T_s (as discussed in Section II) [3]. Thus, we scale the achieved rate by this factor when we derive our results.

Traffic: We consider two types of traffic patterns and run separate simulation experiments to evaluate the performance of our approach with each of these types. First, we use a CBR traffic source with a packet size of 512 bytes. When the end-to-end delay is measured an arrival rate of 2 packets/s is chosen (used previously in studies on the performance of routing protocols [10]). Whereas for measuring the throughput the sources send data packets as fast as they can. This way we can measure how much data can the network deliver under saturated conditions. Second, we use a HTTP server at the application layer to generate web traffic. The reason for choosing HTTP is that, as shown by comprehensive measurement studies [12], most of the traffic generated by wireless users today is due to web access. For both the traffic types considered, we vary the number of source-destination pairs to vary the overall load. The source and destination nodes are chosen randomly from among the nodes in the network and remain active for the entire simulation time. All the simulations are run for 500 seconds of simulation time.

Mobility: In order to model mobility, we use two types of models. First, we use the random waypoint model. In light of the work in [43] the velocity is chosen randomly between 1 and 19 m/s in order to avoid nodes from degenerating to static behavior. Since, this might still have the nodes converge to slow speeds, we also carry out simulations wherein nodes move in randomly chosen directions but with constant speed. We vary this constant speed; the results provide a better evaluation of the changes in the performance of our approach with variations in speed. With either model, when a node reaches its randomly chosen destination point, it invokes a pause time of 25 seconds before it chooses a new destination point in the region of interest and moves in the direction of the chosen point.

Topology: We divide space into square units, each of which is 250 m in length and breadth. This corresponds to the nominal transmission range of a wireless card with a SISO link complying with the IEEE 802.11 standards. We consider a region that is 9 units by 9 units in area. We deploy 200 nodes in this area.

¹²For repeatability of the experiments, the default OPNET parameters are used.

¹³Note that instead of extending the range, one could instead attempt to achieve higher data rates over shorter ranges with virtual MISO or MIMO links. The study of such an alternative approach is beyond the scope of this work.

Comparisons: We compare the performance of our multi-layer approach to that with the *standard stack* using SISO links. The “standard stack” employs the IEEE 802.11 at the MAC layer and the DSR protocol at the routing layer. We also use an idealized approach wherein the shortest virtual MISO path is identified; we compare the routes computed by our approach with that computed with this approach.

Protocol parameters and settings: For the IEEE 802.11 MAC protocol and DSR, we use the standard parameter values as these are set in the official implementation of these protocols in OPNET [13]. However, for the virtual MISO links, the time-outs are extended to account for the extra delay introduced by multiple control message transmissions (local and the virtual MISO based transmissions) and the pilot tones. For example, after the M-RTS is transmitted, the time-out for receiving the M-CTS is the value of the standard time-out for a (SISO) CTS plus: (a) the transmission time for a local CTS (on a SISO link by the receiver) and, (b) the time taken for the pilot tone transmissions (by the receiver and its collaborators). The time for pilot tone transmissions (together) is set to $32\mu\text{s}$ in our simulations; this period is typically considered to be sufficient for channel estimation [21] and includes the time-intervals between pilot tone transmissions. As discussed in Section II, we assume the deployment of either TR-STC, ST-OFDM or a decision feedback equalizer to cope with synchronization issues. When an anycast is invoked, the time-out value is increased by an additional time of $\text{Anycast}_{\text{wait}}$ of $k \times \text{SIFS}$ seconds, to account for the possible back-offs (due to contention among the receivers that attempt to respond) prior to an anycast response. We recall that when a M-RTS message is sent in an anycast mode, the source of the anycast solicits responses only from a restrictive set. This set consists of those nodes on the primary path that are at most k hops away from the source; k is a protocol parameter and is set to 4 in our simulations¹⁴. The nodes in the restrictive set, trigger time-outs at the end of which they would attempt to respond with an M-CTS message. Each member of this set, chooses its time-out in proportion to its hop-count from the source. If the hop-count of the member from the source is j (where $1 \leq j \leq k$), the time-out is set to $(k - j) \times \text{SIFS}$ seconds.

If a virtual MISO link (unicast) fails in facilitating the transmission of a packet within 4 attempts (as in the IEEE 802.11 standard), it is considered to have failed. If a unicast link were to fail, the transmitting node would first resort to an anycast on the virtual MISO link. If this anycast were to fail (after a single attempt), DSR would invoke a new route discovery. The caching methods deployed with DSR are also adopted in our simulation models.


Single hop SISO transmissions: If the destination node is within a single SISO hop of the transmitter, we simply perform a SISO transmission (as per the IEEE 802.11 MAC rules) as opposed to invoking a virtual MISO transmission.

Performance Metrics. We used the following metrics to evaluate the performance of our approach¹⁵:

- 1) **Network Throughput:** We define this to be the number of data bits that are successfully transmitted at the MAC layer by all the nodes in the network

¹⁴As long as k was small, the results were not very sensitive to this parameter.

¹⁵We do not measure physical layer capacity improvements with virtual MISO in this paper. There are previous efforts that report these improvements [36] [35].



(a) Network Throughput.

(b) Average End-to-End delay.

Fig. 8. Static Topologies and CBR Traffic: The performance in term of throughput and average end-to-end delay

- 2) **Average Delay:** This represents the per packet average end-to-end delay.
- 3) **Average number of route failures:** This represents the average number of routes failures per second (considering all of the RERR or route error messages, generated by all the nodes in the network).
- 4) **Average number of hops:** This represents the average length of the constructed routes in the network in terms of logical hop count. A logical hop could either consist of a SISO link or a virtual MISO link.

B. Results and Discussion

Experiments with Static Topologies: For the first set of experiments, we use static topologies (of 200 nodes) and vary the number of source-destination pairs. To begin with, CBR Traffic was used. In Figure 8(a), the throughputs achieved with the multi-layer approach (virtual MISO or VMISO) and the standard stack (SISO) are depicted. We observe an increase in throughput by as much as 100%. The increase here is mainly due to the extended range possible due to the diversity gain; the increase in range results in shorter path lengths and is especially beneficial for UDP applications (such as CBR)¹⁶. The shorter paths also facilitate a more expeditious delivery of packets to their respective destinations. As shown by the results in Figure 8(b), the average end-to-end delay is decreased significantly in comparison to that with the standard stack (by up to 50 %). We note that with 60 flows, the improvements become less pronounced due to effects of congestion; however, with our approach, we still achieve improvements of the order of 50 % in terms of throughput, over what is achieved with the standard stack. Note that this delay reduction is in spite of the additional time taken for the transmission of a packet at the MAC layer (due to two transmissions per packet, i.e. local and cooperative).

Performance in Mobile Scenarios: Next, we consider mobile scenarios; the mobility models that we use were described earlier. For the simulations with constant speeds, we have conducted experiments over a range of speeds. However, we only present the results of our experiments with speeds of 15 m/s and 20 m/s; we also present results wherein we deploy the generic random waypoint model with and represent it with (1,19) m/s in the depicted plots.

The first set of results that we present are again with CBR traffic. In Figure 9(a) we depict the throughputs achieved with

¹⁶We corroborate this claim by examining the average path length in terms of the logical hops later.

our multi-layer approach and that with the standard stack. A significant increase in the throughput is observed with the multi-layer approach; the increase becomes more pronounced at higher mobilities. In high mobility, where the nodes move at 15 m/s, an **increase in throughput of about 150%** is observed with our multi-layer approach. The more dramatic improvements (as compared with the static scenarios) are due to the robustness that our multi-layer approach provides to link failures which increase with mobility. As discussed in section III, the use of our virtual-MISO anycast mechanism coupled with the higher coverage reduces the sensitivity of connections to the mobility of the nodes. As shown in Figure 9(c), **the number of route failures with our multi-layer approach is lower by up to 60%** as compared with the standard stack. Our scheme also **decreases the average end-to-end delay as observed in Figure 9(b) (by up to 75 %)**; the reduction in delay is both due to the extended range and the reduction in broadcast periods resulting from a reduced number of route discovery attempts. At extremely high mobility (20 m/s), the ability of our approach to cope with mobility is reduced only slightly; an improvement of about 140 % is still observed in terms of throughput, over what is achieved with the standard stack.

For the next set of results we use HTTP traffic. Ten HTTP flows that represent client/server connections are randomly initiated in our topology of 200 nodes. The application layer parameters (arrival rate, packet size) are taken from real traces and provided as a part of the OPNET distribution. We perform an exhaustive set of experiments as with CBR traffic; however, for brevity we depict only the throughput. The behavioral results with regard to latency are similar. With HTTP, we again observe a significant improvement in the throughput (by approximately 100%) as shown in Figure 10. We note that the improvement is similar in spirit to that with CBR traffic. Thus, these results demonstrate that the performance benefits are enjoyed irrespective of whether UDP (CBR) or TCP(HTTP) is deployed as the transport layer protocol.

We point out that there is however, a notable difference in the performance results with HTTP as compared with that with CBR traffic. With HTTP, the mobility does not influence the achieved performance gains in the throughput; the throughput is consistently higher by 100% with our approach (as in the static case with CBR traffic). We attribute this phenomenon to the bursty nature of HTTP traffic and the manner in which DSR maintains routes. In more detail, unlike CBR, HTTP consists of bursty traffic, which means that a connection has alternating active and idle periods of data transfer. During the idle periods, the routes maintained in the DSR caches might expire due to the lack of packet arrivals. When a new packet arrival occurs after an idle period, the routes will have to be discovered again. This process is carried out irrespective of whether or not there have been link failures due to mobility in the interim. Thus, the path formed with virtual MISO links is also *deemed* an expired path and will have to be rediscovered with HTTP traffic. When CBR was used, due to sustained data packet arrivals, DSR caches remained valid unless there was a link failure. Upon link failure, the virtual MISO anycasts significantly helped in route recovery as discussed earlier. This benefit is reduced drastically with HTTP traffic.

Path dilation in comparison to an ideal routing policy: Our routing approach was built on top of a SISO based approach to ensure backward compatibility with SISO based networks; in




Fig. 9. Mobility and CBR Traffic: The performance in term of all the defined metrics




Fig. 10. Network Throughput with Mobility and HTTP traffic.

addition this design feature provided a robustness to mobility by facilitating anycasting at the MAC layer. Due to this design policy, it is possible that our protocol may miss shorter paths that could potentially be available if only virtual MISO transmissions were used. To deal with this issue, our protocol employs a dynamic route configuration mechanism which was described previously in Section III. However, there are still pathological cases wherein our approach can generate longer paths than what is possible in an idealized setting. Consider the topology in Fig. 11. For making the figure easier to follow, we show only the relay nodes that constitute the path and their cooperative neighbors. Suppose that node A has a packet to transmit to node H. The primary path discovered by DSR using SISO links is A,B,C,D,E,F,G,H. Our protocol creates virtual MISO links and constructs the shorter virtual MISO path via A,D,G,H. However, an omniscient shortest-path routing protocol would discover the path A,I,H. The reason is that node I is within the coverage range of a virtual MISO link from node A while it is unreachable if SISO is used.

In order to quantify the possible path dilation (as compared to the omniscient scheme) from our approach, we perform simulations. We also compare our approach with the SISO based approach to quantify the reduction in path length as compared to this approach. We use the same simulation settings as in the experiments of subsection B (static randomly generated topology). For this experiment, 10 random topologies are generated and 30 source destination pairs are randomly chosen. We compute the average path length over all the topologies and for each protocol. We list the computed values in Table I. The results demonstrate that our approach compares well with the omniscient scheme (the path lengths are only longer by about a logical hop). Furthermore, the improvements as compared to the SISO based approach are dramatic; the logical hop count is reduced to about half of that with the SISO approach.

Effect of Node Density on the Throughput: By its very de-




Fig. 11. MISO vs. VMISO. If real MISO are employed, node A will communicate with node H, over the path A,I,H. On the other hand, if virtual MISO links are used, the path will be A,D,G,H.

	Average Path Length
SISO	6.12
VMISO	2.44
Omniscient	2.1

TABLE I
AVERAGE ROUTE LENGTH IN NUMBER OF HOPS.

sign, cooperation is better possible with increased density. Thus, our last objective is to investigate the impact of the node density on the performance of our protocol and understand as the regime in which the performance benefits are best achieved. For these experiments, we use the same settings as in the scenarios with static topologies and CBR traffic; however, we vary the node density. The metric of interest is the network throughput. The results are depicted in Table II. As we see, when the average node degree is less than four¹⁷, which is the number of cooperating neighbors required by our implementation, the performance of our protocol reduces to that of SISO. However, as the node density increases, and therefore there are enough neighbors for the virtual MISO links to be built, the throughput increases significantly. Nevertheless, after the node degree is more than 13.48, the throughput does not increase any further. The reason is that, after a certain density threshold, the probability that the required number of neighbors (5) are available in order to create the virtual MISO links, does not change significantly.

V. RELATED WORK

In this section, we describe previous related work. We first discuss previous efforts on the use of cooperative transmissions

¹⁷We assume that every node is aware of the number of one hop neighborhood for this experiment. If a node realizes that it does not have the sufficient number of neighbors to create a virtual MISO link, it switches to a SISO transmission and follows the IEEE 802.11 rules.

Number of Nodes	SISO Throughput	VMISO Throughput	Average Node Degree
80	381164	381164	3.37
200	1030620	2254714	8.16
350	1056864	2509960	13.48
450	996853	2577981	16.7

TABLE II
NETWORK THROUGHPUT FOR VARIOUS NODE DENSITIES.

and receptions (or the formation/use of virtual MISO or MIMO links); these are primarily physical layer related efforts. Later, we describe related work on the use of antenna arrays in ad hoc networks in brief; most of the work here has been on the use of directional antennas.

Use of Cooperative Transmissions and Receptions: The benefits of virtual MIMO systems at the physical layer have been recently studied in [36], [34], [35], [17], [30], [9]. In [34], the authors demonstrate how the coverage (cell size) can be increased for establishing the reverse link in cellular systems with the use of two cooperating mobile transmitters, instead of one. This work also presents an information theoretic model that demonstrates that cooperating transmitters can improve the link level throughput. The work in [35] considers a practical CDMA based implementation of cooperative transmissions at the physical layer and develops receiver architectures and detection algorithms to enable the reception of such transmissions. In our description of virtual MISO links, we have used simple replication of information by the cooperating transmitters. The use of error control codes in order to represent information (instead of simply repeating information) was suggested by [17]; in particular, this work suggests the use of punctured codes.

The benefits of using virtual antenna arrays (referred to as cooperative diversity in these efforts) from a theoretical perspective have also been shown by Laneman [18]. In particular, Laneman shows analytically that, with two cooperating nodes, full diversity (i.e., diversity of order 2) can be achieved. This implies that the outage probability decays in proportion to the inverse of the square of the signal-to-noise-ratio ($1/SNR^2$) with cooperative diversity rather than $1/SNR$ which is the rate of decay without cooperative diversity.

In light of this fact, several studies have recently explored the use of cooperative diversity as a potential tool to improve the power efficiency of wireless communications [9], [26], [19]. In [9], the authors study the energy efficiency achieved with actual and virtual MIMO systems. It is shown that the energy savings (with respect to SISO systems) with virtual MISO systems increase linearly with the distance between the transmitter and the receiver. In [19], the authors provide first indications that cross-layer design considering cooperative diversity may result in significant energy efficiency. In particular, the authors suggest graph theoretical methods for the selection of paths in a randomly constructed network employing cooperative transmissions. Neither [9] or [19], however, propose practical cross-layer approaches that tie the physical, MAC and network layer mechanisms together, for use with virtual antenna arrays.

In addition to the creation of virtual antenna arrays, other types of *multi-user* diversity have also been studied. With accumulative broadcasts, proposed by Maric and Yates, a receiver utilizes multiple partially recovered (SNR below the decoding

level) receptions of the same packet [26]; the authors show that this can improve the lifetime of the network. Furthermore, by use of *temporal diversity*, the receiver may combine partial information received at different times [32] (hitch-hiking), and this may be used to achieve energy efficient broadcasting [1].

While the authors in [30] recognize the importance of multiple-access and higher layer issues, there have been very few efforts, to the best of our knowledge, that have considered the importance of interactions between layers and the implementation of appropriate protocols to exploit cooperative communications. Recently, Barros and Servetto [6] have introduced the reachback problem in sensor networks. The sensor networks are designed so that multiple sensor nodes in the field report back their measurements to a collection center. In practice, the center may be located far away from the sensor nodes in the field; this may prevent individual nodes from establishing a connection with the center on their own. In order to alleviate this problem, the authors propose the use of *cooperative* transmissions of multiple sensor nodes. In their work the authors assume that the nodes have the knowledge of the entire field of the observed samples; this requirement may entail large overheads and could be difficult in practice. Furthermore, the authors do not address MAC or routing issues that arise due to cooperative transmissions.

There have been a lot of efforts on enabling the phase synchronization of cooperative transmissions at the receiver [28], [5], [29], [24], [27], [16], [23], [41], [25]. These solutions were discussed earlier in Section II and thus, we do not repeat the discussion here. Other efforts on the synchronization of transmissions for cooperative broadcasts are [15], [14]. In [15] and [14], the authors propose time-synchronization methods that facilitate cooperative transmissions of a large number of sensor nodes to a single collection center.

Use of multiple antennas in ad hoc networks: There has been some work on the use of specialized antennas in ad hoc networks. Most of the work though assume the use of steerable or directional antennas wherein the antenna can focus energy in a desired direction. Examples of such efforts may be found in [33],[39],[8],[4] and,[20]. In [33], the author points out that the requirement in terms of the size of antenna arrays could make physical deployment on mobile nodes difficult on the spectral bands in use today. In particular, the size would be a factor for operations on the 900 MHz or the 2.4 GHz bands.

Recently, there has been some work on the use of MIMO links in ad hoc networks by Sundaresan and Sivakumar [38], [37]. These works assume that separate flows are established between the different antenna elements of the sender and receivers to yield a spatial multiplexing gain. The joint use of the antenna elements to provide robustness to fading effects has not been considered. Furthermore, the authors assume that antenna arrays are physically mounted on mobile devices.

VI. CONCLUSIONS

In this paper, we design and develop a novel, multi-layer approach for distributed spatio-temporal communications (deployment of virtual MISO links) in mobile ad hoc networks. The use of virtual MISO links is motivated by the size constraints that limit the deployment possibilities of physical antenna arrays. In order to establish a virtual MISO link, groups of nodes simultaneously transmit appropriately encoded information to a single

receiver; this provides the performance advantages of an antenna array. A significant advantage of using virtual antenna arrays is that it does not require any additional hardware except for the provision of space-time signal processing capabilities.

Our work complements prior physical layer work on the use of virtual antenna arrays via the construction of novel higher layer protocols to support virtual MISO deployments. The proposed approach requires a synergistic and tight collaboration between the physical, medium access control, and routing layers. The key physical layer property that we exploit is an increased transmission range achieved due to the diversity gain enabled by collaborative transmissions on a virtual MISO link.

The results from our simulations demonstrate the success of our approach in terms of facilitating the use of and exploiting virtual MISO links. Our approach provides: (a) a significant improvement in the end-to-end performance in terms of throughput and delay, and (b) robustness to link failures due to mobility and interference. More specifically, with our approach, we achieve up to a 150% increase in the end-to-end throughput (in mobile scenarios) and a 75% decrease in the end-to-end delay. Our results also demonstrate a reduction in mobility induced route failures by as much as 60%.

As possible future research directions, we plan to investigate more sophisticated selections of cooperating nodes that, for example, could help minimize the energy consumption in the network or increase network longevity. The characterization of the performance of virtual MISO when used to increase data rates instead of the transmission range, is another future direction that we will explore. Finally, the use of virtual MISO links for enabling MAC layer broadcasts will also be considered in our future efforts.

REFERENCES

- [1] M. Agarwal, J.H. Cho, L. Gao, and J. Wu. Energy efficient broadcast in wireless ad hoc networks with hitch-hiking. In *IEEE INFOCOM*, 2004.
- [2] S. Alamouti. A simple transmit diversity technique for wireless communications. *IEEE Journal on Select Areas in Communications*, Vol. 16, No. 8, pp. 1451-1458, October 1998.
- [3] A.Paulraj, R.Nabar, and D.Gore. Introduction to space-time wireless communications. Cambridge Univ. Press, May 2003.
- [4] L. Bao and J. Garcia-Luna-Aceves. Transmission scheduling in ad hoc networks with directional antennas, 2002. Proceedings of ACM MOBIHOC.
- [5] G. Barriac, R. Mudumbai, and U. Madhow. Distributed beamforming for information transfer in sensor networks. In *Proceedings of IPSN'04*, 2004.
- [6] J. Barros and S. D. Servetto. Network information flow with correlated sources. Submitted to the IEEE Transactions on Information Theory, 2003.
- [7] J. K. Cavers. An analysis of pilot symbol assisted modulation for rayleigh fading channels. *IEEE Trans. Veh. Technol.*, 40:686-693, 1991.
- [8] Romit Roy Choudhury, Xue Yang, Nitin H. Vaidya, and Ram Ramanathan. Using directional antennas for medium access control in ad hoc networks. In *MOBICOM*, 2002.
- [9] S. Cui, A. J. Goldsmith, and A. Bahai. Energy-efficiency of mimo and cooperative mimo techniques in sensor networks. *IEEE Journal on Selected Areas of Communications*, Vol. 22, No. 6, August, 2004.
- [10] Samir Ranjan Das, Charles E. Perkins, and Elizabeth E. Royer. Performance comparison of two on-demand routing protocols for ad hoc networks. In *INFOCOM (1)*, pages 3-12, 2000.
- [11] S. Haykin and M. Moher. Modern wireless communications. Pearson Education, Upper Saddle River, NJ, 2005.
- [12] Tristan Henderson, David Kotz, and Ilya Abyzov. The changing usage of a mature campus-wide wireless network. In *Proceedings of ACM MOBIHOC*, 2004.
- [13] <http://www.opnet.com>.
- [14] A. Hu and S. D. Servetto. Algorithmic aspects of the time synchronization problem in large-scale sensor networks. *ACM/Kluwer Journal on Mobile Networks and Applications (MONET)*, Special Issue on Wireless Sensor Networks. To appear.
- [15] A. Hu and S. D. Servetto. Asymptotically optimal time synchronization in dense sensor networks. In *Proceedings of ACM WSNA '03*, 2003.
- [16] D. Richard Brown III, Gregory B. Prince, and John A. McNeill. A method for carrier frequency and phase synchronization of two autonomous cooperative transmitters. In *IEEE SPAWC*, 2005.
- [17] M. Jananai, A. Hedayat, T.E. Hunter, and A. Nosratinia. Coded cooperation in wireless communications: Space-time transmission and iterative decoding. *IEEE Trans. Signal Processing*, 52:362-370, 2004.
- [18] J.N.Laneman. Cooperative diversity in wireless networks: Algorithms and architectures. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, August 2002.
- [19] A. Khandani, J. Abounadi, E. Modiano, and L. Zhang. Cooperative routing in wireless networks. In *Allerton Conference on Communications, Control and Computing*, October 2003.
- [20] Thanasis Korakis, Gentian Jakllari, and Leandros Tassioulas. A mac protocol for full exploitation of directional antennas in ad-hoc wireless networks. In *Proceedings of ACM MOBIHOC '03*, 2003.
- [21] Srikanth V. Krishnamurthy, Anthony S. Acampora, and Michele Zorzi. Polling-based media access protocols for use with smart adaptive array antennas. *IEEE/ACM Trans. Netw.*, 9(2), 2001.
- [22] E. G. Larsson, P. Stoica, E. Lindskog, and J. Li. Space-time block coding for frequency selective channels. In *IEEE ICASSP*, 2002.
- [23] Xiaohua (Edward) Li. Space-Time Coded Multi-Transmission Among Distributed Transmitters without Perfect Synchronization. *IEEE Signal Processing Letters*, 11(12), December 2004.
- [24] Xiaohua (Edward) Li, Mo Chen, and Wenyu Liu. Application of STBC-encoded cooperative transmissions in wireless sensor networks. *IEEE Signal Processing Letters*, 12(2), Feb 2005.
- [25] E. Lindskog and A. Paulraj. A transmit diversity scheme for channels with intersymbol interference. In *IEEE ICC*, 2000.
- [26] I. Maric and R. D. Yates. Cooperative multicast for maximum network lifetime. *IEEE JSAC*, Special Issue on Wireless Ad Hoc Networks, June 2004.
- [27] Yan Mei, Yingbo Hua, Ananthram Swami, and Babak Daneshrad. Combating synchronization errors in cooperative relays. In *IEEE ICASSP*, 2005.
- [28] Raghuraman Mudumbai, Joao Hespanha, Upamanyu Madhow, and Gwen Barriac. Scalable feedback control for distributed beamforming in sensor networks. In *IEEE ISIT*, 2005.
- [29] F. Ng, J.-H. Hwu, M. Chen, and X. Li. Asynchronous space-time cooperative communications in sensor and robotic networks. In *IEEE ICMA*, 2005.
- [30] A. Nosratinia, T.E. Hunter, and A. Hedayat. Cooperative communication in wireless networks. *IEEE Communications Magazine*, 2004.
- [31] K. Pahlavan and P. Krishnamurthy. Principles of wireless networks: A unified approach. Prentice Hall PTR, 2002.
- [32] John G. Proakis. Digital communications. 4th ed. McGraw Hill, 2001.
- [33] Ram Ramanathan. On the performance of ad hoc networks with beamforming antennas. In *Proceedings of ACM MOBIHOC*, 2001.
- [34] A. Sendonaris, E. Erkip, and B. Aazhang. User Cooperation Diversity-Part I: System Description. *IEEE Transactions on Communications*, 51(11), November 2003.
- [35] A. Sendonaris, E. Erkip, and B. Aazhang. User Cooperation Diversity-Part II: Implementation Aspects and Performance Analysis. *IEEE Transactions on Communications*, 51(11), November 2003.
- [36] Andrew Sendonaris, Elza Erkip, and Behnaam Aazhang. Increasing uplink capacity via user cooperation diversity. In *IEEE International Symposium on Information Theory (ISIT)*, 1998.
- [37] Karthikeyan Sundaresan and Raghupathy Sivakumar. A unified mac layer framework for ad-hoc networks with smart antennas. In *MobiHoc*, pages 244-255, 2004.
- [38] Karthikeyan Sundaresan, Raghupathy Sivakumar, and Mary Ann Ingram. A fair medium access control protocol for ad-hoc networks with mimo links. In *INFOCOM*, 2004.
- [39] Mineo Takai, Jay Martin, Rajive Bagrodia, and Aifeng Ren. Directional virtual carrier sensing for directional antennas in mobile ad hoc networks. In *Proceedings of ACM MOBIHOC*, 2002.
- [40] Vahid Tarokh, Hamid Jafarkhani, and A. Robert Calderbank. Space-time block codes from orthogonal designs. *IEEE Transactions on Information Theory*, 45(5):1456-1467, 1999.
- [41] Shuqngqung Wei, Dennis Goeckel, and Matthew Valenti. Asynchronous cooperative diversity. In *CISS*, 2004.
- [42] A. Wittneben. A new bandwidth efficient transmit antenna modulation diversity scheme for linear digital modulation. In *IEEE ICC*, 1993.
- [43] Jungkeun Yoon, Mingyan Liu, and Brian Noble. Random waypoint considered harmful. In *INFOCOM*, 2003.