An Experimental Case for SIMO Random Access in Multi-antenna Multi-hop Wireless Networks

Ahmed Khattab
The Center for Advanced Computer Studies (CACS), University of Louisiana at Lafayette, Lafayette, LA 70504
akhattab@ieee.org

***Abstract*—In this paper, we demonstrate that multiple concurrent *asynchronous* and *uncoordinated* Single-Input Multiple-Output (SIMO) transmissions can successfully take place even though the respective receivers do not explicitly null out interfering signals. Thus motivated, we propose simple modifications to the widely deployed IEEE 802.11 MAC to enable multiple non-spatially-isolated SIMO sender-receiver pairs to share the medium. Namely, we propose to increase the physical carrier sensing threshold, disable virtual carrier sensing, and enable message in message packet detection. We use experiments to show that while increasing the peak transmission rate, spatial multiplexing schemes such as employed by the IEEE 802.11n are highly non-robust to asynchronous and uncoordinated interferers. In contrast, we show that the proposed multi-flow SIMO MAC scheme alleviates the severe unfairness resulting from uncoordinated transmissions in 802.11 multi-hop networks.**

I. Introduction

Random access Medium Access Control (MAC) protocols are susceptible to packet losses and unfairness in throughput distribution when the competing senders are not able to coordinate their transmissions. Increasing the underlying physical layer rate increases the rate of *only* MAC contention-winning flows. However, flows which are not able to win medium access still suffer very low throughput, despite their potentially high physical layer rate as was experimentally shown in [1] using commodity IEEE 802.11n hardware.

A key reason for the failure of the IEEE 802.11n to provide fairness is that the multi-antenna physical layer is used *only* to increase the per-link throughput (assuming fading and receiver noise are the only sources of randomness). However, such a physical layer does not counter *uncoordinated* or *hidden* senders, but instead relies on the 802.11 MAC to prevent their negative effects. In multi-hop random access networks, wherein senders are not necessarily within range, asynchronous interference from concurrent transmissions is unavoidable since nodes necessarily take uncoordinated transmission actions (e.g., starting time, power, or rate, etc.).

In this paper, we consider the additional source of randomness at the physical layer resulting from random and unpredictable interference from *uncoordinated* transmissions, and ask what is the best use of multiple antennas in this case. Unlike related information-theoretic [2]–[4] and MIMO MAC [5]–[8] approaches, we consider the uncoordinated interference scenario wherein a flow is unable to infer the interferers' channels, rates, power selections, or starting times. The contributions of the paper are as follows.

First, we experimentally demonstrate that in the presence of uncoordinated interference, Single-Input Multiple-Output (SIMO) receive diversity provides increased robustness for a wide range of signal to interference ratio (SIR) compared to MIMO spatial multiplexing schemes. Such schemes require significantly high SIR margin at the receiver to attain the promised gains. Thus, SIMO robust transmission is suitable when senders do not *apriori* know the SIR at the receivers. For example, using the WARP platform [9], we show that a 4×4 MIMO flow has to be more than 15 dB above an uncoordinated interferer to attain the $4\times$ throughput gain. Meanwhile, a 1×4 SIMO flow's transmission is almost error-free at -5 dB SIR. Furthermore, we show that for a given cumulative interference, SIMO reliability worsens with fewer uncoordinated interferers. Therefore, allowing more SIMO transmissions to take place at low transmission power is less harmful to the ongoing transmissions than few high-power ones.

Second, we show how simple modifications to the IEEE 802.11n protocol can exploit the reliability of SIMO links to alleviate the consequences of uncoordinated transmissions. The main idea is to allow for *multiple* asynchronous spatially-proximate SIMO transmissions to take place rather than attempting to ensure that a single spatially-isolated flow uses MIMO to increase its rate (as the case with 802.11n). We show that this can be achieved by suitably increasing the physical carrier sensing threshold and disabling virtual carrier sensing.[1] Consequently, a sender can initiate a new transmission – even if other nearby transmissions are currently taking place – as long as the cumulative interference in its vicinity implies a sufficient interference margin for an additional SIMO transmission. Meanwhile, a receiver must be able to lock on to a new arriving packet after receiving the preamble of an unintended packet. Thus, enabling the Message-In-Message (MIM) [10] 802.11 feature is mandatory for our multi-flow MAC approach. Our results show that our SIMO MAC alleviates the unfairness of legacy MIMO MAC in problematic topologies.

The paper is organized as follows. We motivate the SIMO case in Section II. In Section III, we experimentally demonstrate SIMO robustness to uncoordinated interference. Then, we present simple modifications to 802.11 to exploit SIMO flows to alleviate unfairness in Section IV. We discuss related work in Section V and conclude in Section VI.

II. System Model and Motivation

In this section, we define the system model and the uncoordinated transmissions' problems in random access networks.

A. System Model

We consider a multi-hop ad hoc network wherein individual sender-receiver pairs may not be mutually within transmission range of each other. Each node is equipped with a half-duplex transceiver with $N > 1$ antennas tuned to the same frequency.

[1] Virtual carrier sensing prohibits a node from transmitting after an overheard 802.11 packet header is decoded until after the packet's duration field indicates that the transmission will be completed.

All nodes in the network are competing for a common channel access according to 802.11-like MAC based on carrier sense multiple access (CSMA). Competing senders independently decide whether to transmit or not based on the measured interference energy currently on the shared medium. RTS/CTS collision avoidance mechanism is disabled to overcome the associated overhead.[2] Nodes are not globally synchronized.[3] Transmissions from out of range senders *(i)* asynchronously start at arbitrary time instants, and consequently, *(ii)* their channel information cannot be estimated by the receiver. Hence, explicit interference cancellation techniques are not applicable. We refer to such out of range transmissions as ***uncoordinated interference***. A receiver estimates the channel information of *only* its respective sender to retrieve the transmitted information using a maximal ratio combining architecture. Once a receiver locks to the preamble of its *intended* packet, it remains in the receive mode until the end of the packet reception.

B. Problem Definition and Motivation

The 802.11n standard exploits a multi-antenna physical layer to realize high link rates via MIMO spatial multiplexing. However, the MAC layer still attempts to ensure a single spatially isolated sender-receiver pair. In single-cell networks wherein all nodes are within range, the 802.11n MAC largely ensures such behavior. Consequently, in such idealized scenarios without hidden terminals, a transmission can utilize all of the MIMO channel's degrees of freedom (DoF: defined as the minimum of the link's transmit and receive antennas [13]) to increase its data rate via spatial multiplexing. Unfortunately, in multi-hop networks (or networks in which all nodes are *not* mutually within range yielding hidden nodes), the 802.11n MAC cannot ensure such interference-free condition due to the existence of uncoordinated senders. The resulting collisions have been shown not only degrade the channel utilization but also lead to severe unfairness in the throughput received by competing sender-receiver pairs [11].

Motivation. Our objective is to exploit the underlying multi-antenna physical layer to counter uncoordinated out of range interfering senders in 802.11 multi-hop networks. Instead of allowing only a single contention-winning flow per area to exclusively use MIMO to increase its rate, we ask what is the best multi-antenna scheme that allows multiple flows to simultaneously transmit. Unlike prior work [5]–[8] which necessitate coordination among interfering flows, and hence require network-wide synchronization, we consider uncoordinated random access system model in which senders independently take their transmission actions and receivers are unaware of the interference channel information.

III. Experimental Evaluation of SIMO Robustness to Uncoordinated Interference

In this section, we make the case for SIMO random access networks by experimentally demonstrating that *(i)* SIMO links provide increased robustness to uncoordinated (synchronous and asynchronous) interference as compared to MIMO links with spatial multiplexing for realistic channel conditions, *(ii)* the robustness of SIMO links degrades with fewer high-power interferers as compared to many low-power interferers with the same cumulative interference power. Thus, enabling more SIMO flows to concurrently transmit (even at low-power) results in highly reliable links even with the receivers do not explicitly handle the consequent uncoordinated interference. These properties are generalizations of Telatar's conjuncture in the seminal paper [14] (assuming synchronized transmissions).

A. Testbed Setup

We use the WARPLab framework which uses the WARP FPGA boards as wireless interfaces while the MIMO processing is performed by a PC to which up to 16 WARP boards can be connected via an Ethernet switch. Figure 1(a) depicts a schematic of an example 3-node WARPLab setup. The PC acts as a centralized controller which is used to *(i)* generate the transmitted data streams, *(ii)* realize the MAC scenarios under investigation (e.g., it controls which senders transmit, when, and with how many antennas), and *(iii)* post-process the received data streams. We implement a V-BLAST MIMO spatial multiplexing transceiver with maximal ratio combining receive architecture. Each antenna is used for the transmission of an independent packet with power P. Thus, a $n \times N$ link implies n times the SIMO link transmission power. This transceiver automatically downscales to a SIMO one when only one antenna is used for transmission. The reported results are the average of five measurements.

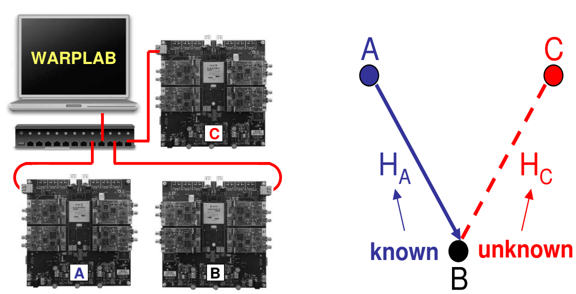

(a) Schematic of a 3-node WARPLab setup. (b) Single interferer scenario.

Fig. 1. A schematic of the 3-node baseline experiment.

Two precautions are made to ensure that any given MIMO link can utilize all the available spatial DoF. First, the inter-spacing between the antennas connected to a WARP board is greater than twice the operating wavelength. Second, the experiments are performed in an indoor propagation environment characterized by rich scattering with none-line-of-sight paths between any sender and any receiver. For all experiments, individual link measurements indicated full rank channel matrices for all links. Therefore, all our findings are attributed to the interaction between multiple transmissions rather than ill-conditioned links.

Tested MAC Approaches. We consider two MAC scenarios of the two uncoordinated senders: *(i)* a random access MAC scenario in which the interfering senders are unaware of A's ongoing transmission. Therefore, such interferers transmit during A's packet transmission. This MAC scenario models hidden senders in 802.11n multi-hop networks. *(ii)* A prefect scheduler MAC scenario in which the PC allows only one sender to transmit at a time with equal channel access time guaranteed per sender. The perfect scheduler does not include coordination and information distribution overhead of an actual protocol and serves as a benchmark for comparison.

[2]While RTS/CTS can mitigate the effect of hidden terminals in some cases, it has no benefit in other cases and increases overhead in all cases [11].

[3]Closed-loop MIMO schemes (e.g., beamforming, stream control, or optimal antenna selection) necessitate global node synchronization and information exchange overhead that outweighs the MIMO throughput gain [12].

B. Reliability of SIMO Links with Uncoordinated and Asynchronous Interference

In this experiment, we show that SIMO robustness to uncoordinated interference compared to MIMO spatial multiplexing holds for real systems *irrespective of the temporal alignment of the uncoordinated interference* and *without explicit interference cancelation*. Figure 1(b) depicts the experiment scenario. For the tagged link *AB*, we consider a single equidistant interfering node *C* using the same number of antennas. We independently tested the link from each sender to the common receiver to ensure that all senders obtain identical goodput in the absence of interference. Node *B* only estimates channel H_A and does not have the channel information of the interfering sender. We consider two cases for uncoordinated random access: A full- and mid-packet interference scenarios wherein *C*'s transmissions start exactly with and exactly in the middle of A's transmissions, respectively.

SIMO Robustness. Figure 2 depicts the goodput of link *AB* normalized to the SIMO goodput for both MAC approaches in 4-antenna systems for full- and mid-packet interference. $N \times N$ MIMO spatial multiplexing requires high SIR to achieve the $N\times$ goodput gain, regardless the temporal alignment of the uncoordinated interference. For example, link *AB* needs to be more than 15 dB above the interferer to obtain the goodput gain for full-packet interference. In contrast, SIMO $1 \times N$ transmissions result in very low outage probability, and thus achieves almost unity normalized rate even at low SIR values. The reliability of a SIMO link increases with the number of receive antennas due to the increase in the available DoF. Thus, a 1×4 transmission achieves unity normalized goodput at -5 dB while a 1×2 requires at least 2.5 dB SIR. We omit the 2-antenna system results for space limitations.

Note that the goodput of the 1×4 link *AB* is only 55% of its upper bound. The perfect scheduler goodput bound can still be attained in the presence of *C*'s interference if *A* and *C* use two antennas each. Yet, *A* must be at least 5 dB stronger than *C*. Therefore, an optimal MAC must track the SIR at the receiver side in order to choose the optimal spatial multiplexing degree. However, such coordination *(i)* induces overhead which further degrades the goodput, *(ii)* does not guarantee optimality as the SIR can change during the course of a packet transmission due to the asynchronicity of uncoordinated transmissions. Interested readers are referred to [1] for further discussion of SINR-based selection of the spatial multiplexing degree.

Temporal Alignment of Interference. For a given goodput, the SIR required given mid-packet interference is lower as compared to full-packet interference – despite similar general trends. This is attributed receiver *B*'s ability to obtain an accurate estimate of H_A compared to the full-packet interference case wherein *A*'s preamble is interfered by *C*'s transmission. Moreover, only half *A*'s payload bytes are susceptible to interference from *C*. The packet delivery ratio increases with the decrease in the overlap time between the two transmissions. We do not further investigate this issue.

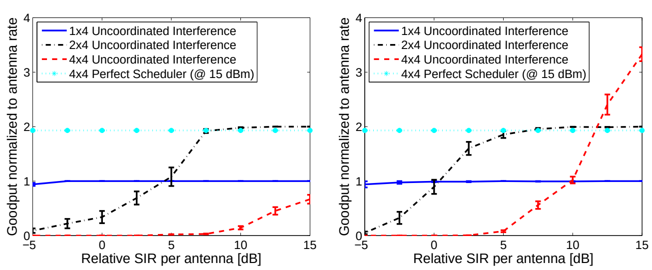

(a) Full-packet interference. (b) Mid-packet interference.

Fig. 2. Goodput of link *AB* versus the relative SIR. For the perfect scheduler, *A* is interference-free and transmitting with power equal to 15 dBm

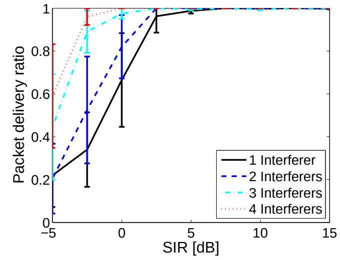

Fig. 3. Goodput of link *AB* versus the absolute SIR.

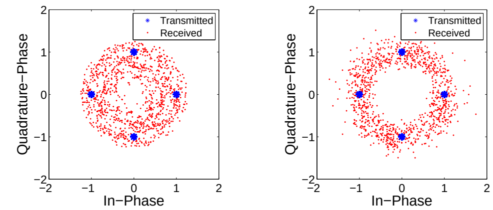

(a) Single Interferer. (b) Four Interferers.

Fig. 4. Transmitted and received symbol constellations of 10 packets at the matched filter output for 0 dB absolute SIR for the 2-antenna system.

C. Multiple SIMO Transmissions Scenarios

In this experiment, we demonstrate that the effect of uncoordinated interference is weakened with the increase in the number of the interferers for a fixed cumulative interference.[4] We plot the packet delivery ratio of the tagged link *AB* given a fixed cumulative interference power while varying the number of interferers for a 2-antenna receiver in Figure 3. All senders (including *A*) are at the same distance of from the receiver *B*. All interferers start at the same time. The x-axis represents the absolute SIR (defined as the ratio of sender *A*'s power and the sum of the power of all interferers). We split the cumulative power equally among all interferers.

As shown in Figure 3, the packet delivery ratio increases as the number of interferers increases. This is attributed to the fact that the assumed single-user detection treats the cumulative interference as a single source of randomness, and does not have any information to null out individual interferers. The complex-valued signals transmitted from different interferers undergo different complex-valued channels. The cumulative interference vector is the phasor addition of the received interference vectors. Due to the randomness in the phases of the received signals form different interferers, their phasor addition yields cumulative interference of amplitude much less than the sum of the amplitudes of the individual received interference vectors. In other words, multiple weak interferers will have the same effect of a single interferer with same total power only if their received vectors are in phase (which is true with probability

[4]We also studied the case in which the relative flow-wise SIR is fixed and showed that a $1 \times N$ flow achieves 70-83% delivery ratio in the presence of $N-1$ interferers each with 0 dB relative SIR. As the flow-wise SIR increases, a SIMO flow achieves higher delivery ratio (or tolerate more interferers).

zero). Figure 4 illustrates the decrease in the variance of the received symbols with the increase of the number of interferers due to the reduction of interference effect given multiple interferers. This implies that a few strong uncoordinated interferers are more harmful to an ongoing transmission compared to many weak interferes. Similar performance was obtained when we repeated this experiment for a SISO system.

IV. SIMO Fair Random Access MAC

In this Section, we present the physical and MAC layer modifications required to allow multiple non-spatially-isolated SIMO flows to successfully communicate and demonstrate the gains compared to legacy MIMO MAC.

A. 802.11n Physical and MAC Layer Modifications

Our multi-flow SIMO MAC relies on allowing sender-receiver pairs to initiate new transmissions – if possible – even if proximate sender-receiver pairs are active. However, the 802.11n standard targets having a single transmission per unit area. Only two modifications are required to 802.11n physical and MAC layers to enable multi-flow SIMO communications. Namely, *(i)* preventing a single sender-receiver pair from exclusively reserving the medium by suitably increasing the physical carrier sensing threshold and disabling virtual carrier sensing, and *(ii)* modifying the packet detection at the physical layer to track the relative changes in the RSSI value to allow receivers to lock to new packets in the presence of active transmissions.

1) Exclusive Medium Reservation Prevention: According to the 802.11 MAC, a sender decides whether or not to transmit based on physical and virtual carrier sensing. If the physical carrier sensing implies that the cumulative energy on the medium exceeds a certain threshold over a DIFS period, a sender infers the existence of a nearby transmission and defers. Legacy physical carrier sensing threshold is chosen to allow a single transmission per carrier sensing area. In contrast, we propose to suitably increase the carrier sensing threshold in order to allow for multiple transmissions to coexist per carrier sensing area. Therefore, a sender can still transmit - even if other nearby transmissions are currently taking place - as long as the cumulative interference is below the new threshold. Furthermore, a sender must continually perform physical carrier sensing in order to initiate its transmission once the cumulative interference implies non-saturated medium. Hence, virtual carrier sensing must be disabled such that a sender does not have to wait for a particular ongoing transmission to end.

2) Message-In-Message (MIM) Packet Detection: In order for multiple asynchronous flows to successfully take place, the receivers of different flows must be able to lock to the preamble of a new packet if the current active transmissions are not of interest (i.e., not intended to this particular receiver). This can be easily achieved be monitoring the RSSI changes after receiving the preamble of an unintended packet. If the relative change in RSSI exceeds a certain threshold, a receiver infers the existence of a new packet and locks on to the new preamble. This functionality can be easily implemented by enabling the 802.11 optional Message-In-Message (MIM) feature [10]. MIM is an enhancement to 802.11 physical layer to address physical layer capture by locking to a new arriving stronger packet. However, MIM is not implemented in all 802.11 chipsets as it is not mandated by the standard. For our multi-flow MAC approach, the MIM feature (or similar packet detection mechanism) is mandatory.

B. Performance Evaluation

Here, we demonstrate the performance gains of our SIMO MAC. We study the fairness and the goodput achieved by the modified 802.11 on top of a SIMO physical layer compared to legacy 802.11 running on top of a MIMO $N \times N$ spatial multiplexing physical layer. Our fairness index is the sum of the log utility which is maximized via the proportionally fair flow rates [15]. We developed a custom event driven simulator using MATLAB. We extract the physical layer parameters of the WARPLab platform and use our experimental data to build a lookup table that is used by the simulator for learning of the physical behavior. For both the SIMO and MIMO MAC realizations, we use DIFS, SIFS, and mini-slot durations of 36, 20, and 20 μs, respectively. Binary exponential backoff is used for contention windows in the range [31, 1023]. The maximum retry limit is 7 attempts before a packet is discarded.

Information Asymmetry Topology. According to legacy MIMO 802.11, flow *Aa* receives significantly low goodput compared to flow *Bb* in the information asymmetry topology depicted in Figure 5(a). While *A*'s packets collide at its intended receiver *a*, *B*'s packets are successfully received at *b*. Normalizing the goodput of each flow by that of a single $N \times N$ flow, the total log utility is -1.19 and -1.22 for the 2- and 4-antenna systems, respectively. The proportionally fair maximum log utility is -0.6 achieved when both flows obtain half the $N \times N$ flow rate. In contrast, the proposed SIMO MAC approach allows a to successfully receive most of its intended packets despite *B*'s uncoordinated transmissions. Thus, the goodput of flow *Aa* increases, and hence, the total log utility increases to -0.84 and -1.14 for the 2- and 4-antenna systems, respectively. While SIMO MAC log utility is close to optimal in the 2-antenna case, the log utility degradation is due to SIMO low per-flow goodput in the 4-antenna system.

Flow in the Middle Topology. In the topology depicted in Figure 5(b), the senders of the two out of range MIMO 802.11 flows *Aa* and *Bb* independently and successfully transmit to their respective receivers while the sender of the middle flow is deferring as long as either or both flows are active. Hence, flow *Cc* receives almost zero goodput. The total log utility for the 2- and 4- antenna systems is -4.67 and -4.62, respectively; which is much lower than the proportionally fair utility of -0.83 achieved by a normalized goodput of 0.67, 0.33, and 0.67. In contrast, the proposed SIMO approach allows the middle flow to transmit despite the ongoing activities of the outer flows. In the 4- and 2-antenna systems, the achieved goodput is (0.27, 0.26, 0.27) and (0.46, 0.09, 0.46), yielding log utility of -1.7 and -1.69, respectively. In the 2-antenna system, the goodput degradation of the middle flow is because such flow is receiving interference from two strong interferer. Recall that the performance of our SIMO approach significantly dependens on the instantaneous interference unlike legacy 802.11 which targets a binary on-off interference model. For example, if the inter-flow distance slightly increases to $1.2d$ while keeping the sender-receiver distance at d, the log utility increases to -1.2 as the achieved rates increase to 0.5, 0.21, and 0.5 due to reduced interference.

Hidden Senders Scenarios. Finally, we consider uncoordinated senders scenarios wherein long-term fairness is not the main concern (i.e., all flows operate under identical interference conditions). We consider 2- and 4-flow scenarios wherein all senders are hidden from each other and all senders are at the

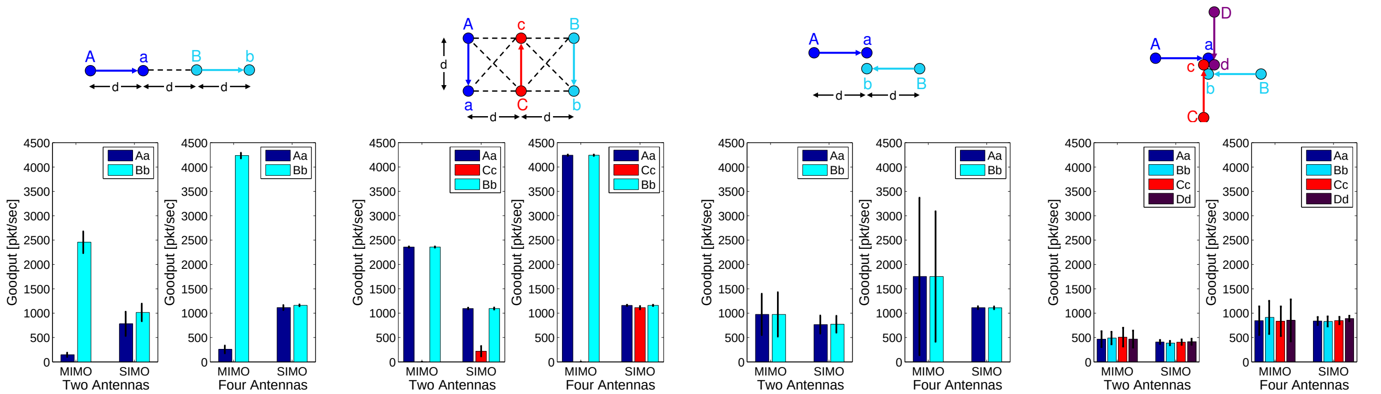

(a) Information asymmetry. (b) Flow in the middle. (c) 2 Hidden senders w/ collocated receivers. (d) 4 Hidden senders w/ collocated receivers.

Fig. 5. Example problematic multi-hop topologies with uncoordinated transmissions. Solid arrows indicate data flows while dashed lines imply nodes are within transmission range. Senders *A* and *B* are uncoordinated (out of range) senders in topologies (a)-(c). All senders in topology (d) are out of range.

same distance from the collocated receivers. Both the SIMO MAC and MIMO MAC approaches yield almost equal throughput per flow as depicted in Figures 5(c) and (d). However, SIMO MAC results in less collisions, and hence, better short-term fairness as indicated by the variance of the goodput depicted via the error bars. In the 2-flow topology depicted in Figure 5(c), the goodput achieved by the SIMO MAC is less than the MIMO MAC in the 4-antenna case due to SIMO low per-flow rate. However, as the number of competing flows increase, the number of simultaneously active flows increase with the number of antennas, and the SIMO average flow rate approaches the MIMO flow rate (both almost equal to 0.2). Even for the 2-antenna system, the per flow goodput in the 4 flow topology (0.17) is close to the MIMO case (0.19).

V. Related Work

MIMO-specific MAC. The upcoming 802.11n MIMO standard suffers from the same severe unfairness and starvation problems encountered in single antenna networks [1]. In contrast, [5]–[8] address fairness by allowing multiple simultaneous transmissions to coexist in the same channel using joint MIMO signaling techniques such as stream control, beamforming, or optimal antenna selection. However, protocols employing these mechanisms result in an overhead due to network synchronization and channel acquisition that significantly degrades the system throughput as was empirically shown in [12]. MIMO multiple transmissions in asynchronous networks was addressed only in [1]. with a scheme that cannot be incrementally deployed to the IEEE 802.11n.

Alternative Random MAC Approaches. Due to the inefficiency and unfairness problems of random access in single radio multi-hop networks, several alternative approaches have been proposed in the literature. Examples include the use of multiple orthogonal channels and/or multiple radios, and the use of directional antennas. Even though such alternatives have the potential of addressing the hidden terminals problem and its consequences, additional resources are required to realize interference-free parallel transmission and new medium access challenges are introduced (e.g., coordination and deafness problems). Alternatively, [16] presents new physical layer design, employing joint decoding of collided packets, to counter hidden terminals. In contrast, we address random access inefficiencies via simple MAC modifications in single-channel multi-hop networks with conventional multi-antenna physical layer.

VI. Conclusions

In this paper, we propose using the multi-antenna physical layer to increase the number of spatially proximate transmissions instead of increasing the per-link rate. A sender uses a single antenna for transmission while a receiver uses all of its antennas to provide diversity gain. We experimentally demonstrate that such a physical layer is robust to uncoordinated and asynchronous interfering transmissions. We present the simple modifications required to the 802.11 protocol to enable multiple SIMO flows to concurrently share the medium. We show that such a multi-flow SIMO MAC alleviates the severe unfairness of legacy 802.11 MIMO MAC in multi-hop topologies.

References

[1] A. Khattab, et al., "Asynchronous randomized allocation multi-antenna medium access in multi-hop networks," in *Proc. of IEEE ICCCN*, 2008.

[2] R. Blum, "MIMO capacity with interference," *IEEE Journal on Selected Areas in Communications*, vol. 21, no. 5, pp. 793–801, Jun. 2003.

[3] N. Jindal, et al., "Rethinking MIMO for wireless networks: Linear throughput increases with multiple receive antennas," in *Proc. of IEEE ICC*, 2009.

[4] A. M. Hunter, et al., "Capacity scaling of ad hoc networks with spatial diversity," in *Proc. of IEEE ISIT*, 2007.

[5] K. Sundaresan, et al., "A fair medium access control protocol for ad-hoc networks with MIMO links," in *Proc. of IEEE INFOCOM*, 2004.

[6] P. Casari, et al., "DSMA: an access method for MIMO ad hoc networks based on distributed scheduling," in *Proc. of ACM IWCMC*, 2006.

[7] R. Bhatia and L. Li, "Throughput optimization of wireless mesh networks with MIMO links," in *Proc. of IEEE INFOCOM*, 2007.

[8] M. Park, et al., "Cross-layer MAC design for wireless networks using MIMO," in *Proc. of IEEE Globecom*, 2005.

[9] A. Khattab, et al., "WARP: A flexible platform for clean-slate wireless medium access protocol design," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 12, no. 1, Jan. 2008.

[10] N. Santhapuri, et al., "Message in message (MIM): A case for reordering transmissions in wireless networks," in *Proc. of ACM HotNets*, 2008.

[11] M. Garetto, et al., "Modeling per-flow throughput and capturing starvation in CSMA multi-hop wireless networks," *IEEE/ACM Transactions on Networking*, vol. 16, no. 4, Aug. 2008.

[12] S. Gaur, et al., "Interfering MIMO links with stream control and optimal antenna selection," in *Proc. of IEEE Globecom*, 2004.

[13] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.

[14] I. E. Telatar, "Capacity of multi-antenna gaussian channels," *European Transactions on Telecommunications*, vol. 10, no. 6, pp. 585–595, 1999.

[15] F. Kelly, A. Maulloo, and D. Tan, "The rate control for communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 409, pp. 237–252, 1998.

[16] S. Gollakota and D. Katabi, "Zigzag decoding: Combating hidden terminals in wireless networks," in *Proc. of the ACM SIGCOMM*, 2008.