

Fair Randomized Antenna Allocation in Asynchronous MIMO Multi-hop Networks

Ahmed Khattab, Ashutosh Sabharwal and Edward W. Knightly
ECE Department, Rice University, Houston, TX 77005
{khattab, ashu, knightly}@ece.rice.edu

Abstract—Previous work has shown that CSMA-based protocols such as the IEEE 802.11 can yield flow starvation in multi-hop wireless networks. While prior protocol designs alleviated such starvation by utilizing MIMO, they require global synchronization to temporally align transmissions and require per-packet distribution of the complete channel state information. In this paper, we experimentally show that MIMO networks based on pre-802.11n medium access, the state-of-the-art for asynchronous MIMO CSMA, *worsen* starvation as compared to SISO networks. Consequently, we design an *asynchronous* MIMO MAC protocol that counters starvation. We show that randomized and non-greedy antenna allocation coupled with local residual capacity estimation results in previously-starving nodes capturing a fair share of system resources while simultaneously exploiting throughput gains available to multi-antenna systems.

I. INTRODUCTION

The upcoming IEEE 802.11n MIMO standard [1] promises performance gains compared to SISO systems by utilizing spatial diversity (increasing link reliability) or spatial multiplexing (increasing link capacity) [2]. However, because the MIMO physical layer employs CSMA/CA at the MAC layer, the 802.11n standard and its variants will suffer from the same severe unfairness and starvation problems encountered in single antenna networks [3]. Indeed, we show experimentally using pre-802.11n devices, that MIMO *worsens* unfairness in key starvation scenarios. Such CSMA/CA starvation is attributed to the asymmetric and incomplete views of the wireless channel for contending flows in multi-hop networks [3]. Use of MIMO simply to improve the performance of individual links provides a further advantage to the winning flows.

Multiple transmit and receive antennas can also be used for beamforming (also called stream control [4]) and interference cancellation, respectively. Thus, in contrast to the IEEE 802.11 standards, multiple simultaneous transmissions can coexist in the same channel. However, protocols employing these mechanisms, such as [5]–[10], require network-wide synchronization and channel information of all interfering transmitters at each receiver in order to null out their signals. While such synchronous MAC protocols address fairness by allowing multiple simultaneous transmissions, the overhead due to network synchronization and channel acquisition significantly degrades the system throughput as was empirically shown in [11].

In this paper, we develop a framework and protocol to (i) enable simultaneous and asynchronous transmissions in

multi-antenna multi-hop networks and (ii) employ randomized and non-greedy resource allocation to counter starvation. Our framework defines a class of asynchronous random access protocols subsuming MIMO systems via the following two components:

Residual Capacity Estimation: The first part of the framework is an interference measurement mechanism that estimates the current network conditions at both the transmitter and receiver. The residual capacity estimate is based on two key ideas: (i) determination of the transmitter and receiver’s residual SINR margin via multi-bit quantization of the sensed interference. This contrasts with carrier sensing, in which a binary transmission decision is made based on whether the sender’s measured interference energy level is greater than a threshold. (ii) Because transmitters and receivers have different channel states, they share their residual capacity estimation information to make a joint decision. Hence, the transmission decision is appropriated for the available channel resources at both link endpoints irrespective of the surrounding topology.

Randomized Resource Allocation: The second part of the framework is a *randomized* decision making policy that uses the measured channel conditions to perform a randomized mapping of measurements into node resource allocations. Random actions counter inherent inaccuracies in the channel condition estimations (e.g., binary exponential backoff and information hiding [12]). We apply randomization to the dimension of resource allocation not only to counter random estimation errors, but also to alleviate flow starvation. In particular, we map a discretized transmitter-receiver measurement to a set of potential resource allocations. For example, a particular discretized sender-receiver measurement pair could map to a transmit antenna set such as use of {2, 3, or 4} antennas; likewise, if the allocated resource is transmit power, the measurement can map to a set of transmit power levels. The transmitter-receiver pair *randomly* selects a value within this set according to a pre-defined distribution.

The policies that map residual capacity estimation regions to resource usage sets and assign probabilities to each member of the set yield vastly different protocol behavior. For example, we will show that excluding the maximum number of antennas from all resource allocation sets yields a non-greedy policy that, coupled with random allocation of the number of transmit antennas, ensures that no flow consistently and greedily captures the channel, possibly starving other flows.

To realize the above framework in multi-antenna networks,

we present ARAM, the **Asynchronous Randomized Allocation Multi-antenna** medium access protocol. The proposed antenna allocation mechanism maps the residual capacity estimate of the channel conditions to a randomized choice of the number of transmit antennas. The main objective of our randomized mapping is a non-greedy policy that yields a sufficient SINR margin for previously-starving flows to initiate *simultaneous* communication if needed. Thus, nodes that transmit do not use all available resources for all time, thereby leaving an opportunity for other flows to also transmit. Consequently, ARAM contrasts with existing protocols such as 802.11n that allow a single flow to repeatedly capture the medium, starving other flows in some topologies. We note that because we fix the transmit power per antenna and vary the number of antennas, this policy has the effect of potentially varying the transmit power per transmission. Therefore, our framework can be applied to single antenna systems employing power- and/or rate-control. However, joint consideration of power, rate, and antenna allocation is beyond the scope of this work.

Finally, we develop a randomized allocation policy for ARAM that maximizes the average transmission rate subject to employing a non-greedy randomized allocation strategy. To define this policy, we compute the achievable rates and outage probabilities for a transmission under asynchronous, multi-antenna, interferers. We use simulations to evaluate a number of ARAM policies compared against 802.11n-like systems.

The remainder of this paper is organized as follows. In Section II we describe the network model and pre-802.11n measurements. In Section III, we present the randomized resource allocation framework and devise ARAM. We study interference in asynchronous MIMO networks and calculate antenna allocation probabilities in Section IV. In Section V, we present simulation results, and we conclude in Section VI.

II. PRELIMINARIES AND 11N MEASUREMENTS

A. System Model

We consider an asynchronous random access network in which each node is equipped with a single half-duplex wireless interface with $N > 1$ antennas. We consider single-hop flows in a multi-hop network topology (i.e., a topology in which all nodes are not mutually within range). We define the MIMO transceiver as follows: The transmitter of flow i uses a subset n_i of its N antennas, with each antenna used for the transmission of a single independent stream or packet. For the rest of the paper, we use the terms antenna and stream interchangeably. The transmitted power per antenna (or packet), P_{ant} , and the stream rate, r , are fixed throughout the network, regardless the number of used antennas.

When acting as a receiver, a node uses all of its antennas. The N replicas of a transmitted stream are combined such that the received signal to interference plus noise ratio (SINR) at the output of the combiner is maximized. The receiver learns *only* about the communication channel of its own transmitter via a training sequence for two purposes: (i) to null-out self interference from other streams transmitted from the same

sender and (ii) to select the best antenna(s) to be used for data transmission. A feedback path is available to convey the selection information back to the transmitter. Such feedback overhead between a sender-receiver pair is significantly less than the overhead of estimating the channel information of *all* interferers and feeding it back to the transmitter for beamforming [11].

The aforementioned system model outlines the underlying multi-antenna physical layer. The design of a sophisticated interference-limited MIMO transceiver (e.g., using more advanced diversity receiver architectures) is beyond the scope of this paper. Indeed, the optimal MIMO physical layer design is not known in the case of random access *and* asynchronous networks in which nodes have only incomplete information about inter-node channels in the network. Most information theoretic analysis for MIMO networks, e.g., [13], [14], assumes availability of significantly more channel information than we allow. For example, synchronized MIMO MAC protocols [5]–[10] allow for a dedicated time slot, as shown in Figure 1(a), in which only one sender sends a training sequence at a time so that receivers acquire the channel information of all interfering flows. This allows such protocols to eliminate the received interference.

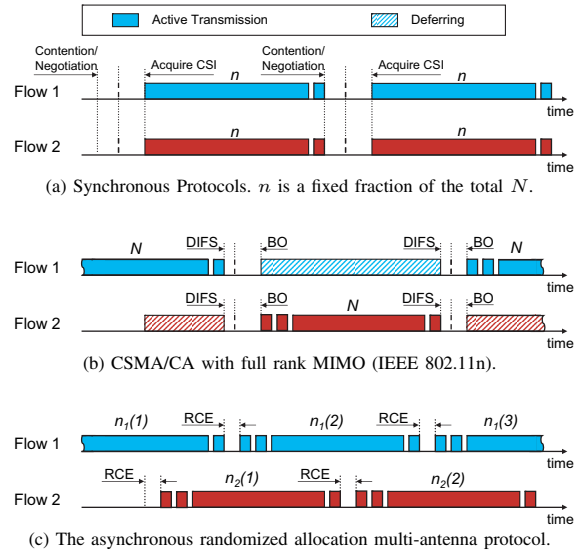


Fig. 1. Time line of the activities of 2 flows under different MAC approaches.

B. Channel Model

According to our network model, the received signal vector at the receiver of flow i is given by

$$\mathbf{y}_i = \mathbf{H}_i \mathbf{x}_i + \sum_{k=1}^{L-1} \mathbf{H}_k \mathbf{x}_k + \mathbf{n} \quad (1)$$

where \mathbf{H}_i is the $N \times n_i$ fading channel matrix between the transmitter of flow i and flow i receiver, and \mathbf{x}_i is the $n_i \times 1$ signal sent by the transmitter of flow i . The $N \times 1$ vector \mathbf{n} represents the additive white Gaussian noise with covariance matrix $E[\mathbf{n}\mathbf{n}^\dagger] = \sigma^2 \mathbf{I}$, where \dagger denotes the conjugate transpose,

σ^2 is the noise power, and \mathbf{I} is the $N \times N$ identity matrix. Elements of the channel matrices and the noise vector are modeled as i.i.d. zero mean, circularly symmetric, complex Gaussian random variables. L is the total number of transmitters.

C. Limitations of CSMA/CA

The IEEE 802.11 standard family represents the state of the art of asynchronous random access mechanisms. The standard adopts CSMA/CA at the MAC layer and consequently targets (within a clique) having a single flow exclusively captures the medium while other flows defer their transmissions as shown in Figure 1(b). In multi-hop networks, all nodes are not mutually within range of each other, hence, nodes have differing channel states. Such asymmetric channel state can result in a backlogged flow capturing the system resources causing other flow(s) to persistently receive very low throughput [3].




Fig. 2. Problematic single-hop scenarios in CSMA multi-hop networks. Nodes connected with dotted lines are within transmission range.

Consider the multi-hop topology shown in Figure 2(a), node B, the transmitter of flow Bb, knows exactly when to contend for the medium because it overhears the activity of flow Aa. On the other hand, node A, the transmitter of flow Aa, has no information about flow Bb and must blindly contend for the medium. The probability of successful transmission of flow Aa packets is close to zero. Similarly, the middle flow Aa in Figure 2(b) will receive very low throughput compared to the outer flows Bb and Cc. Transmitter A will find the medium busy with high probability due to the uncoordinated transmissions of flows Bb and Cc.

The IEEE 802.11n standard will use MIMO to increase the physical layer rate via spatial multiplexing. Unfortunately, allowing the node that captures the channel to transmit at a higher rate does not alleviate flow starvation because starvation originates from the inability to capture the medium [3].

D. 802.11n Measurements

To demonstrate the existence of starvation in MIMO networks, we design the following experiment.

We utilize four laptops, each equipped with a wireless Belkin card that utilizes the *Ralink* RT2860 and RT2820 chipsets. The cards fully comply with the current IEEE 802.11n draft with backward compatibility with the IEEE 802.11b/g standards. The chipset embodies a 2 transmitter, 3 receiver (2T3R) architecture¹ via on-board dipole antennas with 1 dBi antenna gain. We configured the cards in the 802.11n 40 MHz mode with 802.11b compatibility. We use *iperf* to generate fully backlogged UDP traffic sessions at

¹<http://www.ralinktech.com/ralink/data/RT2800.pdf>

transmitting nodes. We arrange the four nodes to form two contending transmitter-receiver pairs in 2 different indoor topologies: a fully-connected topology in which all nodes are within range of each other, and the information asymmetry topology shown in Figure 2(a), in which the transmitter of one flow is out of range of both the sender and the receiver of the other flow.




Fig. 3. The mean and variance of the flow throughput using pre-802.11n devices in the 2 contention scenarios.

Figure 3 depicts the throughput (averaged over 10 measurements, each of length 120 seconds) of each flow in both setups when the RTS/CTS mechanism is both disabled and enabled. While the two flows fairly share the available bandwidth in the fully-connected scenario, one flow receives 68.34 times the throughput of the other flow in 802.11n networks in the information asymmetry topology when the RTS/CTS handshake is disabled. This throughput ratio drops to 12.14 when we repeat the same experiment with 802.11b cards,² as shown in Figure 4. Thus, MIMO worsens the severity of starvation since the flow which exclusively captures the medium transmits at a higher rate compared to the SISO case. Enabling the RTS/CTS handshake does not alleviate starvation, but rather degrades the throughput due to the transmission of such control packets at the base rate.




Fig. 4. IEEE 802.11n worsens starvation in information asymmetry scenarios.

E. Objective

Our objective is to exploit the antenna selection capability of multi-antenna networks to realize fair medium access among interfering flows in multi-hop networks. Our main idea is to prevent a flow from using all antennas for all time thereby

²We do not use the pre-N Belkin cards as they adopt a 1T3R SIMO architecture if configured in the 801.11b mode.

leaving an SINR margin for other contending flows. Two open questions are (i) how to infer the available capacity margin based on the measured interference in the sender and the receiver’s vicinity, and (ii) how to map these estimates into a non-greedy randomized antenna allocation policy that counters starvation while achieving throughput gains available to multi-antenna systems.

III. FRAMEWORK AND PROTOCOL FOR RANDOMIZED RESOURCE ALLOCATION

As described in Section II, different nodes in multi-hop networks have a different view of the channel state. Hence, our methodology has transmitter-receiver pairs collaborate to estimate the residual capacity in their vicinity via passive measurements, and use this information to choose their resource allocation actions. Moreover, the action is *randomized* to overcome the fact that the channel usage estimate is not accurate, and the action is *non-greedy* to prevent a flow from using all available resources. Consequently, other contending flows have an increased opportunity to *simultaneously* transmit.

In this section, we present a general MAC framework for randomized resource allocation. Then, we use this framework to realize an asynchronous counter-starvation MIMO MAC protocol.

A. Framework

The randomized resource allocation framework is based on the following two principles.

1) *Residual Capacity Estimation*: The first part of the framework is passive estimation of the available channel resources. Each node can listen to the on-going transmission activities in its vicinity and can measure the received accumulative energy of these transmissions. The received signal strength spectrum can be divided into multiple regions via multiple thresholds. Hence, nodes can take an appropriate transmission decision (e.g., a node chooses the appropriate transmission rate, power, or coding scheme in single antenna networks; or the number of transmit antennas in MIMO networks) based on the measured residual SINR margin. This contrasts with a binary decision based on a binary measurement as in protocols based on carrier sense.

In multi-hop networks, the transmitter and receiver of a link have different views of the network resources as they are exposed to different spatial regions (e.g., see the network topologies in Figure 2). Therefore, their individual estimates of the residual capacity can differ. We *jointly* use the transmitter-receiver estimates rather than simply using the transmitter’s or the receiver’s individual estimate. A transmitter-receiver pair subsequently share their multi-bit interference estimates to form a joint estimate of the region’s conditions. An appropriated transmission decision is made based on the joint estimation of the available channel resources at both link endpoints irrespective of the surrounding topology. Hence, topology need not be explicitly known by the protocol.³

³While use of mechanisms such as RTS/CTS also employ a joint sender and receiver decision, prior work does not employ joint residual capacity estimation.

2) *Randomized Resource Allocation*: The second part of the protocol framework is the decision making policy based on the measured capacity margin in the flow’s vicinity. The SINR estimates are based on a finite measurement window. Moreover, longer measurement windows are not necessarily useful since the environment is dynamic and the interference is modulated both by bursty traffic, asynchronous initiation and termination of packet transmissions, and channel fading. We propose to use a *randomized mapping* from measurements to transmission actions to counter such inherent inaccuracies in the measurements. Namely, the randomized mapping handles random estimation errors with randomized allocations.

To implement this randomization in resource allocation, we map the measured residual capacity estimation into a set of transmission actions. Each element in that set (e.g., a certain number of antennas, rate, or power level) is associated with a certain probability. Therefore, nodes do not always take the same action for the same measurement. This ensures that nodes are neither consistently greedy nor consistently starving in their choices.

In contrast, 802.11-like resource allocation strategies are deterministic and do not account for inaccuracies in estimating interference. Namely, if a carrier-sensing 802.11 node infers that the interference is below a specific threshold, it transmits at full power with probability 1, i.e., its strategy is deterministic and greedy. However, carrier sense does not provide a complete picture of the channel and resource allocation based on incomplete information leads to starvation or poor fairness characteristics [3].

Regardless, the idea of using randomization to deal with random noise (due to channel fading, source burstiness or channel errors) is commonly used in both network protocol design and information theory. For example, random backoff is a randomizing node action designed to account for the lack of precise information about other transmitters’ queue state, the number of active nodes in the network, etc. Likewise, information-theoretic analysis of jamming channels relies on randomness to increase network capacity, which is again randomizing a nodes actions to counter random actions of a jammer [12]. In contrast, we apply randomization to the dimension of resource allocation not only to counter random estimation errors but also to alleviate flow starvation.

The key challenge in designing randomized mappings is the choice of the actions and their corresponding probabilities. We propose non-greedy policies such that no flow exclusively captures the available channel resources. The main effect of such a randomized mapping is that the interference observed by other contending flows is lower when reduced resources (number of antennas and/or transmit power) are chosen by the current flow. Consequently, an SINR margin remains for other flows to initiate *simultaneous* communication if needed. Thus, nodes that choose to transmit do not use all available resources for all time, thereby leaving an opportunity for other flows to transmit.

B. ARAM Protocol

Asynchronous Randomized Allocation Multi-antenna (ARAM) is a random access protocol that exploits residual capacity estimation and antenna selection to implement the randomized resource allocation framework in MIMO networks. ARAM utilizes a four-way handshake access mechanism. A transmitter precedes data packet transmission with a *request* packet that contains training sequences (sent from each antenna one at a time) to aid the intended receiver in estimating its channel matrix. The request packet includes the multi-bit quantization of the measured interference at the sender. The receiver responds with the number and identity of the selected antennas in a *grant* packet. Finally, data transmission takes place followed by an acknowledgment. The request, grant and ACK packets are sent over a single antenna.

Unlike legacy RTS/CTS handshaking, request/grant handshaking does not exclusively reserve the medium, as non-greedy randomized resource allocation allows for multiple simultaneous transmissions. Instead, the request and grant packets are used to exchange residual capacity estimations and antenna allocation and selection information between the sender and its receiver. The main components of the ARAM protocol are:

Residual Capacity Estimation (RCE): ARAM employs two received signal strength indicator (RSSI) thresholds: T_1 that corresponds to the minimum detectable power of a single antenna, and $T_N = NT_1$ that represents the minimum power received from N interfering antennas. Note that the value of T_N is analogous to the binary carrier sensing threshold in IEEE 802.11n networks. Before making the decision regarding the number of transmit antennas n_i , the receiver needs to know the available capacity margin at its transmitter's vicinity. Therefore, each transmitter communicates to its intended receiver its own RCE region (whether below T_1 , above T_N or in between T_1 and T_N) via 2 bits in the request packet.

Randomized Antenna Allocation: RCE estimates the amount of interference at both link endpoints. Using the RCE information, the receiver then decides how many antennas the sender should use. If information such as network topology, the number of interfering flows, and the number of used antennas per interfering flow is explicitly known, the ideal number of antennas that could be used for the current transmission can be calculated using standard network utility maximization techniques [15]. Probabilistic antenna allocation overcomes the need for such information in asynchronous distributed networks. The number of used antennas and the corresponding probabilities can be chosen to realize different objectives. We propose the following scheme which targets the highest fair throughput in the system (we defer the analysis of this scheme to Section IV).

With RCE information of both the transmitter and the receiver, we identify 3 operation regions: both RCE(TX) and RCE(RX) are below T_1 ; either RCE(TX) or RCE(RX) are above T_N ; or all other combinations of RCE(TX) and

RCE(RX). In the first region, a flow can use all of its antennas as there is no significant interference at both its endpoints; however, the ARAM receiver decides to allow its transmitter to use all of its N antennas only with a certain probability p_N . Otherwise, the receiver chooses to have the sender use $\hat{N} < N$ antennas with probability $1 - p_N$. \hat{N} and p_N govern the fairness-throughput of the protocol. Small \hat{N} (and/or p_N) gives more SINR margin for other contending flows. On the other hand, small \hat{N} degrades the throughput performance of the system. This probabilistic allocation prevents a flow from greedily dominating the medium irrespective of the topology. The residual capacity margin is the key idea by which the ARAM protocol counters starvation so that other (possibly starving) flows can use this capacity margin to transmit some streams. Then, the ARAM protocol will work to achieve fairness among flows since each flow will adapt its transmission strategy according to the changes in the interference level.

In the second RCE region of strong interference at either the sender or the receiver, transmitting a large number of streams can result in collision with the on-going transmissions. Therefore, the receiver decides either to have the sender use a single antenna with probability $1 - p_0$ or to not let the sender transmit with probability p_0 . In the latter case, the transmitter defers and re-measures the medium after a random backoff duration. Finally, for all other combinations of RCE(TX) and RCE(RX), the receiver chooses to use 1 or N antennas with probabilities p_1 and $1 - p_1$, respectively. The probabilistic allocation approach in the above two ARAM operation regions targets to balance the successful transmission opportunities while not degrading the throughput in interference-prone scenarios. Algorithm 1 depicts the ARAM probabilistic antenna allocation procedure.

Algorithm 1 ARAM Probabilistic Antenna Allocation

1. if RCE(TX) < T_1 and RCE(RX) < T_1
 2. $n_i = \hat{N}$ with probability $1 - p_N$;
 3. $n_i = N$ with probability p_N ;
 4. else if $\max(\text{RCE(TX)}, \text{RCE(RX)}) \geq T_N$
 5. $n_i = 0$ with probability p_0 ;
 6. $n_i = 1$ with probability $1 - p_0$;
 7. else
 8. $n_i = 1$ with probability p_1 ;
 9. $n_i = N$ with probability $1 - p_1$;
-

Single-Link Antenna Selection: After determining the number of antennas to be used, n_i , the receiver selects the best n_i antennas of the transmitter's antennas using the channel information learned via the training symbols in the request packet. In ARAM, the receiver uses an exhaustive search algorithm to select the best transmit antennas. Since typical values of N are relatively small, the computational complexity of exhaustive search is limited. The receiver informs its sender the number and identity of the selected antennas via an antenna mask of length N bits.

Data Transmission/Acknowledgment: Upon receiving the grant packet, the transmitter sends n_i different equal-sized data

packets over the selected antenna(s). This is more efficient than sending a single packet at n_i times the single antenna rate as, in ARAM, the control overhead is shared for multiple packets. Another consequence is that, for a given packet size, data transmission time will be fixed regardless of the number of used antennas. The receiver acknowledges the correct reception of individual packets in order to minimize the retransmission overhead. If no acknowledgement was received, the transmitter runs the binary exponential backoff procedure.

IV. RANDOMIZED ALLOCATION POLICY

Recall that ARAM maps measurements of residual capacity to a resource set, with an element of the set chosen according to a pre-defined probability distribution. In this section, we compute the resource set and the corresponding distribution that maximizes the average rate for two contending transmissions with random interference.

We first calculate the achievable flow rates in asynchronous MIMO networks. Since neither the transmitter nor receiver can have complete channel information of all interfering transmitters, we cannot use the same interference channel analysis of [13] and [14]. Our technique is to use the outage probability analysis to calculate the achievable rates.

A. Achievable Rates and Outage in Asynchronous Networks

According to our network model, transmission i consists of n_i independent data streams, each representing a packet, and each is transmitted with rate r . The probability of successfully transmitting a single stream with rate r is one minus the outage probability. Due to the independence of the streams sent by different antennas, the achievable rate of transmission i is

$$R_i = n_i r (1 - p_{out})^{n_i} \quad (2)$$

The information-theoretic outage probability of a *randomly selected* antenna⁴ is defined to be the probability of the event that the mutual information $I(SINR)$ falls below a specific transmission rate r [16], i.e.,

$$p_{out} = Prob[I(SINR) < r]. \quad (3)$$

For our considered diversity scheme, the mutual information per stream is given by [16]

$$I(SINR) = \frac{1}{N} \log_2(1 + SINR). \quad (4)$$

The received SINR of the n^{th} stream at the output of the maximal ratio combiner is given by

$$SINR = \frac{\sum_{m=1}^N P_i |h_{nm}^i|^2}{\sum_{k=1}^{L-1} \sum_{l=1}^{n_k} \sum_{m=1}^N P_k |h_{lm}^k|^2 + \sigma^2} \quad (5)$$

where P_i is the received signal power from node i per antenna, and P_k is the received interference power from node k per antenna.

⁴Unfortunately, the outage probability for the *best* antenna selected cannot be calculated in asynchronous networks. An antenna identified to be the best at the beginning of a transmission can be the worst antenna when the interference channel changes due to misaligned transmissions.

We define the following three terms to describe the output of the maximal ratio combiner: $\gamma_{nm}^i = |h_{nm}^i|^2$, where h_{nm}^i is the channel fading coefficient between the n^{th} transmit antenna of node i and the m^{th} receive antenna; $\gamma_n^i = \sum_{m=1}^N \gamma_{nm}^i = \sum_{m=1}^N |h_{nm}^i|^2$ as the effective SIMO channel of the n^{th} transmit antenna of link i at the output of the combiner; and $\gamma^k = \sum_{l=1}^{n_k} \gamma_n^k$ as the effective interference from node k at the combiner output. Hence, (5) can be rewritten as

$$SINR = \frac{SNR \gamma_n^i}{\sum_{k=1}^{L-1} INR_k \gamma^k + 1} \quad (6)$$

where SNR and INR_k are the mean signal to noise ratio per antenna and the mean interference to noise ratio per antenna received from node k , respectively. SNR and INR_k depend on σ^2 , P_{ant} , antenna heights, the propagation environment, and the inter-node distances.

A set of concurrent transmissions can be characterized via a set of SNR and INR values. For ease of notation, we denote $\gamma = \gamma_n^i$ and $\tilde{\gamma} = \sum_{k=1}^{L-1} \gamma^k$, and assume equal INR for all interfering transmissions. Substituting (6) and (4) in (3), the outage probability can be expressed as

$$p_{out} = Prob\left[\frac{\gamma}{\tilde{\gamma} + \frac{1}{INR}} < (2^{Nr} - 1) \frac{INR}{SNR}\right] \quad (7)$$

In order to evaluate (7), we need to find the probability distribution functions of γ and $\tilde{\gamma}$. Since we assume the channel fading coefficients are normally distributed, γ_{nm}^i is exponentially distributed. The Chi-squared (χ_{2m}^2) distribution with m degrees of freedom nominally applies to the sum of m i.i.d. exponential random variables. The sum of i.i.d. Chi-squared random variables has a Chi-squared distribution with degrees of freedom equal to the sum of the degrees of freedoms of the i.i.d. random variables. Since the channel fading coefficients are i.i.d, γ and $\tilde{\gamma}$ are Chi-squared distributed with $2N$ and $2\tilde{M} = 2 \sum_{k=1}^{L-1} n_k N$ degrees of freedom, respectively.

Thus, the outage probability in (7) is calculated as

$$p_{out} = \int_0^\infty f_{\tilde{\gamma}}(\tilde{\gamma}) \int_0^{A(\tilde{\gamma} + \frac{1}{INR})} f_{\gamma}(\gamma) d\gamma d\tilde{\gamma} \quad (8)$$

where $A = (2^{Nr} - 1) \frac{INR}{SNR}$. For (8), we first evaluate the inner integral

$$\int_0^{A(\tilde{\gamma} + \frac{1}{INR})} f_{\gamma}(\gamma) d\gamma = \frac{1}{\Gamma(N)} \int_0^{A(\tilde{\gamma} + \frac{1}{INR})} \gamma^{N-1} e^{-\gamma} d\gamma \quad (9)$$

$$= \frac{(N-1)!}{\Gamma(N)} \left[1 - e^{-A(\tilde{\gamma} + \frac{1}{INR})} \sum_{s=0}^{N-1} \frac{(A(\tilde{\gamma} + \frac{1}{INR}))^s}{s!} \right] \quad (10)$$

Then, we calculate the outer integral noting that $\Gamma(N) = (N-1)!$ for integer valued N . We have

$$p_{out} = \frac{1}{\Gamma(\tilde{M})} \left[\int_0^\infty \tilde{\gamma}^{\tilde{M}-1} e^{-\tilde{\gamma}} d\tilde{\gamma} - \sum_{s=0}^{N-1} \frac{A^s}{s!} \times \int_0^\infty \tilde{\gamma}^{\tilde{M}-1} (\tilde{\gamma} + \frac{1}{INR})^s e^{-\tilde{\gamma}} e^{-A(\tilde{\gamma} + \frac{1}{INR})} d\tilde{\gamma} \right] \quad (11)$$

$$= 1 - \frac{e^{-\frac{A}{INR}}}{\Gamma(\tilde{M})} \sum_{s=0}^{N-1} \frac{A^s}{s!} \int_0^\infty \tilde{\gamma}^{\tilde{M}-1} (\tilde{\gamma} + \frac{1}{INR})^s e^{-\tilde{\gamma}(1+A)} d\tilde{\gamma} \quad (12)$$

Unfortunately, the integral in (12) is intractable. However, we can asymptotically evaluate upper and lower bounds of p_{out} for interference-limited and noise-limited networks, respectively. We only report the upper bound as our network is an interference-limited one.

Interference-limited Upper Bound: For interference-limited networks where the interference power is much higher than the noise power (i.e., $INR \gg 1$), the intractable integral in (12) is reduced to a tractable one that equals $(s + \tilde{M} - 1)! / (1 + A)^{s + \tilde{M}}$. If the signal to noise ratio is finite, A goes to infinity for infinite INR . In this weak signal to interference scenario, we have certain outage ($p_{out} \rightarrow 1$ almost surely). On the other hand, if the SNR is in the same order of the INR such that $\frac{SNR}{INR}$ is finite, we have $e^{-\frac{A}{INR}} = 1$ and hence (12) is equal to

$$p_{out}^u = 1 - \frac{1}{\Gamma(\tilde{M})(1+A)^{\tilde{M}}} \sum_{s=0}^{N-1} \left(\frac{A}{1+A} \right)^s \frac{(s + \tilde{M} - 1)!}{s!} \quad (13)$$

Thus, the outage probability, and hence, the achievable rates are functions of the stream rate r , the signal to interference ratio $\frac{SNR}{INR}$, and the number of interfering antennas n_k for all interfering transmitters modeled by \tilde{M} .

Substituting with (13) in (2), the achievable rates can be calculated. It is worth noting that if the set of contending transmissions and their SNR and INR values are known *a priori*, standard network utility maximization can be used to determine the ideal number of antennas n_i per transmission. However, realizing such allocation requires a distributed mechanism to identify the set of contending transmissions, measure their SNR and INR values, and coordinate resource allocation among the transmissions.

B. Distributions for Random Transmission Sets

We next present a mechanism to select the ARAM protocol's parameter values for a random interferer. The parameters consist of p_0 , p_1 , p_N and \hat{N} . Recall that the sender-receiver RCE quantization maps to a resource set of $\{0, 1, \hat{N}, N\}$ antennas. Members of the resource set can have different allocation probabilities in different RCE regions.

The mean rate of transmission i for a random contender, \bar{R}_i , is the weighted sum of the rates achieved using all possible values of n_i for all RCE decision regions, i.e.,

$$\bar{R}_i = \sum_{RCE} R_{n_i|RCE} Prob[RCE] \quad (14)$$

where $Prob[RCE]$ is the probability of a given quantized residual capacity estimation. RCE is based on two factors: the relative node locations and the antenna allocations of other transmitters in the network.

One way to calculate $Prob[RCE]$ is to consider a single random interferer. In this case, the probability of each of the possible RCE events depends on the antenna allocation probability of only one other transmitter. Thus, the mean transmission rate in (14) can be rewritten as

$$\bar{R}_i = (R_{\hat{N}|0} + R_{N|0})p_0 + (R_{1|1} + R_{N|1})(p_1 + (1 - p_0)) + (R_{1|\hat{N}} + R_{N|\hat{N}})(1 - p_N) + R_{1|N}(p_N + (1 - p_1)) \quad (15)$$

where $R_{n_i|n_j}$ is the achievable rate of transmission i using n_i antennas given that the interfering transmission is using n_j antennas calculated using (13) and (2).

Under a unit disk transmission model and randomly placed transmitters and receivers, (15) can be used with the probabilities, p_{top} , of different combinations of interfering pairs of transmissions in a random topology as computed in [17]:

$$R_{avg} = \sum_{top} \frac{(\bar{R}_{Aa|top} + \bar{R}_{Bb|top})}{2} p_{top} \quad (16)$$

For each integer value of $\hat{N} < N$, we calculate p_0 , p_1 , and p_N that maximize (16) such that $0 \leq p_0, p_1, p_N \leq 1$.

V. SIMULATION RESULTS

We develop a discrete-time event-driven network simulator using Matlab using the parameter values used to implement the 802.11 protocol in *ns-2* simulator listed in Table I. Each node is equipped with 4 antennas. We assume backlogged UDP flows with non-empty transmit queues. We present ARAM protocol results for $\hat{N} = 3$ and 2 to illustrate the impact of \hat{N} on the throughput-fairness tradeoff. We calculate the values of p_0 , p_1 , and p_N as shown in Section IV. We also tuned the ARAM probabilities to realize a conservative allocation policy in which a sender uses only a single antenna. We refer to this allocation as the ARAM(PF) policy, as it achieves proportional fairness in all two-flow topologies if network utility maximization is used. Our benchmark is an 802.11n protocol which uses N antennas to increase the data rates. In our 802.11n implementation, we transmit N packets per channel access in order to combat the throughput inefficiency resulting from the transmission of control packets at the base rate. We label this protocol as CSMA/CA(N). T_N in the ARAM RCE is calculated to equal the CSMA/CA(N) carrier sensing threshold corresponding to distance of 250m. These results are the average of 20 simulations of length 30 seconds.

Symmetric Interference Scenarios: Figure 5(a) depicts the throughput distribution of 2 flows with symmetric interference effect on one another. Since both senders are in range, CSMA/CA(N) enables both flows to fairly alternate capturing the channel. On the other hand, the ARAM protocol allows both flows to be simultaneously active each using a certain number of antennas according to the interference received from the other flow. CSMA/CA(N) achieves higher throughput than ARAM for values of \hat{N} less than $N/2$ (since CSMA/CA(N) flows alternate using 0 and N antennas, $N/2$

TABLE I
PHYSICAL AND MAC LAYER PARAMETERS.

Parameter	Value
Node TX power	0.2818 W (24 dBm)
Channel BW	20 MHz
Packet size	1000 bytes
Single antenna rate	2 Mbps
PLCP length	192 bits @ 1 Mbps
(R,G,ACK,Data) header	(20,14,14,28) bytes
(DIFS, SIFS, Mini-Slot)	(25, 16, 20) μ sec
(CW _{min} , CW _{max})	(31,1023)

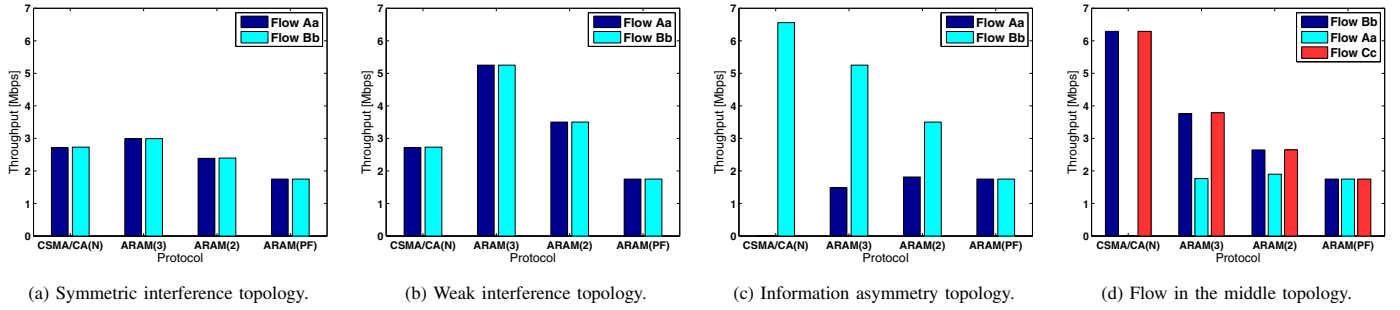


Fig. 5. Flow throughput of example two-flow scenarios under different MAC approaches. In topology (a), $d_{Ab} = d_{aB} = 250m$. In topology (b), $d_{Ab} = d_{aB} = 2d_{AB} = 500m$. In topology (c), $d_{Ab} = 2d_{aB} = 500m$, hence, link Aa receives only 2.7 Kbps in CSMA/CA(N) networks. In topology (d), the distance between a transmitter and receiver of the immediate neighboring flow is 250m. Link Aa receives only 4.7 Kbps with CSMA/CA(N).

antennas are used in average; moreover, the outage probability for CSMA/CA(N) is almost zero). As \tilde{N} exceeds $N/2$, ARAM throughput outperforms CSMA/CA(N) as for ARAM(3).

Interference-Free Scenarios: Next, we consider the case in which the senders of the two flows are in proximity of each other while the interfered receivers are located at far distance (similar results were obtained when sender-receiver roles are switched). Both CSMA/CA(N) and the ARAM(PF) protocol achieve the same throughput as in the symmetric interference scenario as shown in Figure 5(b). In contrast, ARAM(\tilde{N}) yields both flows to alternate using \tilde{N} or N antennas. Hence, it achieves higher throughput compared to CSMA/CA(N).

Asymmetric Interference Scenarios: The ARAM non-greedy policy alleviates the severe unfairness in throughput shares in the problematic information asymmetry scenario shown in Figure 2(a). Note that flow Aa receives only 2.7 Kbps with CSMA/CA(N). However, the superior ARAM fairness comes at the expense of low throughput as shown in Figure 5(c). The choice of \tilde{N} controls the throughput-fairness tradeoff of the ARAM protocol: using $\tilde{N} = 2$ achieves 21.5% lower aggregate throughput than $\tilde{N} = 3$, however, the disadvantaged flow obtains 51% (rather than receiving only 27%) of the advantaged flow throughput. It is the network designer's decision to choose \tilde{N} to reflect the throughput-fairness requirements.

Flow in the Middle Scenario: We use this example network to show that the ARAM protocol capability to prevent flow starvation is not restricted to two-flow topologies. ARAM alleviates the severe unfairness in the throughput of flow Aa as depicted in Figure 5(d). Similarly, \tilde{N} governs the throughput-fairness tradeoff. For the conservative single antenna PF policy, each flow achieves 84% of the theoretical fair throughput in CSMA/CA(N) (one third of the link capacity).

VI. CONCLUSIONS

In this paper, we present a framework for randomized resource allocation in medium access protocols. In this framework, a transmitter-receiver pair jointly estimate the surrounding network conditions and the residual capacity margin. Based on these estimates, a flow probabilistically use its available resources (the number transmit antennas, transmit power, or coding rate). The probabilistic approach counters starvation by not allowing a flow to use all available resources

for all time, thereby leaving a resource margin for other contending flows to initiate simultaneous transmissions. We use this framework to realize the first asynchronous MIMO MAC protocol that alleviates flow starvation.

REFERENCES

- [1] IEEE P802.11 - Task Group N - Meeting update. Available at http://grouper.ieee.org/groups/802/11/Reports/tgn_update.htm.
- [2] A. J. Paulraj, D. A. Gore, R. U. Nabar, and H. Bolcskei, "An overview of MIMO communications - a key to gigabit wireless," *Proceedings of the IEEE*, vol. 92, no. 2, pp. 198–218, Feb. 2004.
- [3] M. Garetto, T. Salonidis, and E. Knightly, "Modeling per-flow throughput and capturing starvation in CSMA multi-hop wireless networks," in *Proceedings of IEEE INFOCOM '06*, Barcelona, Spain, Apr. 2006.
- [4] M. Demirkol and M. Ingram, "Stream control in networks with interfering MIMO links," in *Proceedings of IEEE WCNC*, Mar. 2003.
- [5] P. Casari, M. Levorato, and M. Zorzi, "DSMA: an access method for MIMO ad hoc networks based on distributed scheduling," in *Proceedings of ACM IWCMC*, Vancouver, Canada, Jul. 2006.
- [6] M. Park, S.-H. Choi, and S. M. Nettles, "Cross-layer MAC design for wireless networks using MIMO," in *Proceedings of IEEE Globecom '05*, Dec. 2005.
- [7] M. Park, R. J. Heath, and S. Nettles, "Improving throughput and fairness for MIMO ad hoc networks using antenna selection diversity," in *Proceedings of IEEE Globecom '04*, Dec. 2004.
- [8] K. Sundaresan, R. Sivakumar, M. Ingram, and T.-Y. Chang, "A fair medium access control protocol for ad-hoc networks with MIMO links," in *Proceedings of IEEE INFOCOM '04*, Hong Kong, Mar. 2004.
- [9] R. Bhatia and L. Li, "Throughput optimization of wireless mesh networks with mimo links," in *Proceedings of IEEE INFOCOM '07*, Anchorage, Alaska, May 2007.
- [10] A. Ashtaiwi and H. Hassanein, "Rate splitting mimo-based mac protocol," in *Proceedings of IEEE Conference on Local Computer Networks*, Dublin, Ireland, Oct. 2007.
- [11] S. Gaur, J.-S. Jiang, M. Ingram, and M. Demirkol, "Interfering MIMO links with stream control and optimal antenna selection," in *Proceedings of IEEE Globecom '04*, Dallas, TX, Nov. 2004.
- [12] P. Moulin and J. A. O'Sullivan, "Information-theoretic analysis of information hiding," *IEEE Transactions on Information Theory*, vol. 49, no. 3, 2003.
- [13] S. Ye and R. Blum, "Optimized signaling for MIMO interference systems with feedback," *IEEE Transactions on Signal Processing*, vol. 51, pp. 2839–2848, Nov. 2003.
- [14] R. Blum, "MIMO capacity with interference," *IEEE Journal on Selected Areas in Communications*, vol. 21, pp. 793–801, Jun. 2003.
- [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] J. N. Laneman, "Limiting analysis of outage probabilities for diversity schemes in fading channels," in *Proceedings of IEEE Globecom '03*, San Francisco, CA, Dec. 2003.
- [17] M. Garetto, J. Shi, and E. Knightly, "Modeling media access in embedded two-flow topologies of multi-hop wireless networks," in *Proceedings of ACM Mobicom '05*, Germany, Aug. 2005.