Spoofing Uplink Spatial Multiplexing with Diverse Spectrum

Adriana B. Flores and Edward W. Knightly
Department of Electrical and Computer Engineering, Rice University, Houston, TX

***Abstract*—While 802.11 Access Points (APs) can have as many as eight antennas, form factor constraints often limit mobile clients to a single antenna, thereby precluding spatial multiplexing. In this paper, we design, implement, and experimentally evaluate Diverse Spectrum Spatial Multiplexing (DSSM), the first system to enable uplink spatial multiplexing for clients with a single Wi-Fi antenna. We show how to spoof an un-modified 802.11n-compliant AP for reception of data from a DSSM transmitter that employs diverse spectrum to realize a virtual array with nearby devices. We implement and experimentally evaluate DSSM on a software-defined-radio platform. We demonstrate that DSSM communicates with and spoofs an off-the-shelf MIMO receiver. We demonstrate multiplexing gains that increase PHY throughput from 6.5 Mbps for 1x4 BPSK 1/2 rate coding to 259 Mbps for 4x4 with 64-QAM 5/6.**

I. Introduction

MIMO spatial multiplexing enables transmission of independent data streams using multiple transmit and receive antennas. The multiplexing gains and hence data-rate improvements increase linearly with the *minimum* of the number of transmit or receive antennas. Unfortunately, while 802.11 APs can have as many as eight sub-6-GHz antennas, form factor constraints of mobile devices often limit clients to a single sub 6-GHz antenna. Consequently, the antenna asymmetry between the AP and mobile clients precludes spatial multiplexing for a transmission between a single client and the AP.

In this paper, we design, implement, and experimentally evaluate Diverse Spectrum Spatial Multiplexing (DSSM), the first system to enable uplink spatial multiplexing for clients with a single sub-6-GHz antenna. DSSM enables the receiver (AP) to be compliant with IEEE 802.11n/ac MIMO operating in 2.4/5 GHz, thereby enabling a DSSM client to spoof an unmodified AP to infer that the single-antenna client has an array. In particular, we present the following contributions.

First, we introduce a diverse spectrum architecture that consists of a *short-range network* (typically realized with antenna arrays and codebook based beamforming) and a (relatively) *narrowband long-range legacy network* with clients having a single legacy-band antenna (or in-band antenna) due to the larger wavelength for long range. We use the 60 GHz mm-wave band to exemplify a wideband *short-range network*, and the *legacy network* operates in the 2.4/5 GHz bands. The intended scenario for DSSM is one where the backhaul device (i.e. AP) is only reachable with the legacy band. However, the mobile devices use the mm-wave band to form a short-range network which enables multi-stream transmission with the legacy AP. A DSSM client with uplink data to transmit, termed the Uplink Originator (UO), leverages neighboring clients in the short-range network as Spatial Multiplexer (SM) devices. In this way, the mm-wave band permits integration of distributed devices in a manner that is transparent to the AP. By jointly utilizing high and low frequency bands, the DSSM architecture exploits the speeds of the mm-wave band to realize multiple legacy-band streams, thereby attaining longer range communication of lower frequency bands while also realizing spatial multiplexing for antenna-limited devices.

Second, in order to spoof reception at an unmodified legacy AP, we design a client PHY and MAC layer that coordinates the independent clients to enable uplink spatial multiplexing. The DSSM-PHY introduces three distributed PHY processing blocks that allow a client to operate from SISO to any MIMO configuration despite having a single legacy-band Tx/Rx chain. The DSSM-MAC design ensures clients' 802.11 medium access is unmodified by isolating the UO and the SM clients with the introduction of a multiplexer block in the SM device that acts as a remote partial-replica of the UO-MAC. The spatial multiplexer block ensures that the SM's medium access and protocol stack is unchanged while giving the UO remote control over medium access, retransmission and acknowledgments of its data.

Third, we implement DSSM on a software-defined-radio platform and evaluate its performance using an extensive set of over-the-air (OTA) experiments. DSSM successfully spoofs a commercial off-the-shelf receiver with two single-antenna DSSM transmitters. Moreover, DSSM achieves 16.3 bits/sec/Hz empirical capacity (using estimated OTA channels) for a 4x4 transmission at high SNR, corresponding to 326 Mbps for a 20 MHz channel and 652 Mbps for 40 MHz. Next, we study our system limitations such as time and phase synchronization, as well as stream outage. Through our evaluation we show how resilient DSSM can be to these limitations, for example, DSSM can successfully decode with the packet with the outage of one other distributed transmitter for modulations up to 16-QAM 1/2.

II. Diverse Spectrum Spatial Multiplexing (DSSM)

In this section, we describe the design of DSSM that overcomes the physical bottleneck on the number of streams a *mobile* device can support.

A. Diverse Spectrum Node Architecture

The node architecture of DSSM mobile devices employs dual-band communication: wideband short-range communication and narrowband long-range legacy communication. Wide-band communication achieves high data rates due to GHz-wide channels available at high frequencies (short wavelength). In this paper, we utilize the unlicensed 60 GHz mm-wave band to realize a wide-band network. In contrast, the narrow band employs lower carrier frequency for greater range and penetration. Yet this band is an order of magnitude lower in bandwidth, e.g., 20-80 MHz, and its relatively long wavelength constraints clients to a single antenna. In this paper, we utilize the unlicensed 2.4/5 GHz bands for this purpose and refer to them as the "legacy bands." With both high and low frequency bands, we target to exploit the speeds of the mm-wave band to realize multiple legacy-band streams, thereby attaining both long range and spatial multiplexing in antenna-limited devices.

We introduce two roles for DSSM nodes to enable diverse spectrum spatial multiplexing. First, we denote the client that has uplink data to transmit and seeks to perform spatial multiplexing as the "Uplink Originator" (UO). Second, we denote as "Spatial Multiplexer" (SM) nodes that perform frequency and protocol conversion on the data of the UO to enable spatial multiplexing. A particular station can function as both UO and SM for different transmissions. Note, a SM node is not a relay node, because the data being transmitted by each device (in the virtual array) is different, i.e. each device transmits different symbols of the packet. In contrast to relay transmissions (2 in-band transmissions), DSSM legacy-band multi-stream transmission is simultaneous, the sharing of the data occurs prior to the MIMO transmission in the high-bandwidth out-of-band link not using any of the legacy band resources.

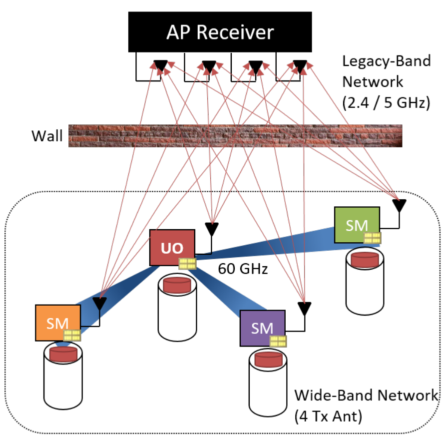

Fig. 1: DSSM network enabling 4x4 uplink spatial multiplexing for single in-band antenna clients.

An example DSSM network is depicted in Fig. 1 with one UO and three SM devices enabling four stream uplink spatial multiplexing and spoofing a legacy AP. The DSSM network is created using the wide-band interconnection that allows devices to be mobile by its wireless nature. Moreover, the high data rates of the wideband short-range network permits communication that is an order of magnitude faster compared to legacy 2.4/5 GHz communication to the AP. While high-rate point-to-point 60 GHz communication is available between devices, we consider that longer-range and non-line-of-sight (NLOS) communication is needed to reach the AP.

B. DSSM Physical Layer

Here we describe the DSSM PHY design for distributed transmitters in legacy bands that enables reception via 802.11n standard APs, i.e., reception via single-user MIMO with spatial multiplexing. DSSM PHY is enabled by the introduction of three distributed PHY processing blocks (shown in Fig. 2) designed to enable MIMO transmissions in devices with a single Tx/Rx chain.

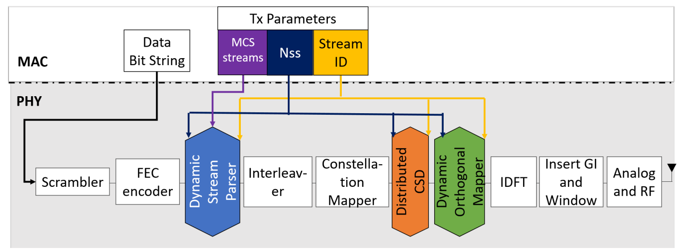

Fig. 2: DSSM PHY design.

Dynamic Distributed Stream Parser. First, we design a Dynamic Distributed Stream Parser (DDSP) such that distributed transmitters mimic the stream parser of a device with multiple co-located radios (viz. an 802.11n multi-antenna transmitter). In contrast to a multi-stream output of a co-located parser, DDSP outputs one stream since DSSM devices only require a single legacy antenna. Moreover, DSSM requires the flexibility to select and adapt the output stream.

To enable the adaptive functionality of the DSSM PHY, we require to input the complete MAC protocol data unit (MPDU) string and the newly designed transmit parameters (also see Fig. 2). Standardized 802.11 transmit parameters include packet length, data rate, transmit power and scrambler initialization. To adapt to the available resources, we additionally include the Number Spatial Streams (N_{ss}), assigned stream ID (i_{ss}) and the MCS for all the streams, as depicted. The MPDU and Tx parameters are generated by the Uplink-Originator-MAC. With this information the DDSP can follow the functionality of the IEEE 802.11n/ac stream parser, retaining bits for selected stream i_{ss} and discarding other bits.

Fig. 3 depicts an example of DDSP for a 4x4 MIMO transmission in which four spatial streams have an equal MCS of 64-QAM with 6 coded bits per single subcarrier (N_{BPSCS}), resulting in 3 bits per stream $s(i_{SS})$ as defined by:

$$s(i_{SS}) = \max\left\{1, \frac{N_{BPSCS}(i_{SS})}{2}\right\}. \quad (1)$$

The output of the parser $x_k^{i_{ss}}$ is a function of the input y_i, as shown for the case of equal MCS by:

$$\begin{aligned} x_k^{i_{ss}} &= y_i \\ i &= (i_{ss}-1)\cdot s + s\cdot N_{ss}\cdot \left\lfloor \frac{k}{s} \right\rfloor + k \bmod s \end{aligned} \quad (2)$$

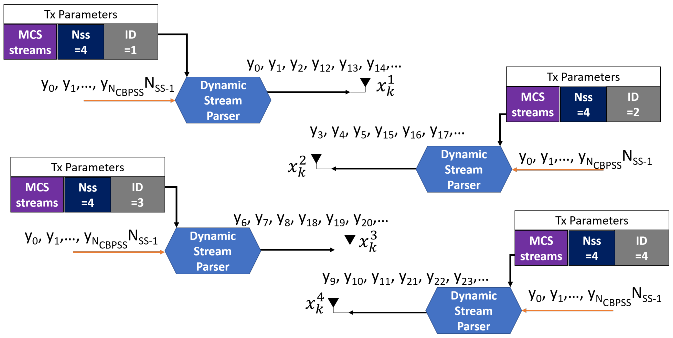

Fig. 3: Four Stream Dynamic Distributed Stream Parser.

where the stream ID (i_{ss}) range is $1 \leq i_{ss} \leq N_{ss}$, s is the number of bits per stream, k is the output bit index ($k = 0, 1, ..., N_{CBPSS} - 1$) and N_{CBPSS} is the number of coded bits per spatial stream defined by the modulation order.

DDSP can dynamically output a spatial stream for any MIMO configuration using Equation (2) and the knowledge of three values provided by the new transmit parameters: Total Number Spatial Streams (N_{ss}), assigned stream ID (i_{ss}) and the MCS for all the streams.

Distributed Cyclic Shift Delay. Cyclic Shift Delays (CSD) permit multi-antenna systems to transmit training symbols through all antennas without causing correlated signals and undesirable beamforming effects. IEEE 802.11n defines a set of CSD values to decorrelate signals for devices with multiple co-located antennas that are spaced a minimum of half wavelength apart. Contrary to co-located antennas, the physical separation of DSSM's distributed transmitters aids signal de-correlation reducing the need for CSD, as DSSM transmitters are expected to be separated by multiple wavelengths. However to ensure training signals are not correlated and to maintain standard compatibility, we apply CSD to all transmit streams as specified in the 802.11n standard.

We introduce the Distributed CSD block, shown in Fig. 2, for distributed transmitters to independently apply Cyclic Shift Delays. Contrary to the standard CSD processing where each CSD block applies a pre-defined CSD value, DSSM's distributed CSD can apply any CSD value of a given MIMO configuration and stream ID assignment. We design the distributed CSD block to implement the complete standard-defined CSD matrix but enable the ability to lookup the CSD value to apply based on the number of spatial streams (N_{SS}) and stream ID (i_{ss}), which are provided by the transmit vector.

Dynamic and Distributed Orthogonal Mapping. DSSM directly maps streams to RF transmit chains because only one spatial stream and one transmit chain is available at each independent transmitter. Even though the complete MPDU is accessible at the distributed transmitters, stream combining is not available unless the MPDU is processed multiple times by the single transmit block to generate all the streams for combining.

With direct mapping, DSSM devices perform omni-directional transmissions, thus we rely on precise *receiver* channel estimation to decode the multiplexed data streams. The DSSM stations simultaneously transmit a set of training symbols, high-throughput long training field (HT-LTF) in 802.11, that allow channel-state-information to be estimated at the receiver (CSIR) through the introduction of the Dynamic Orthogonal Mapping (DOM) Matrix.

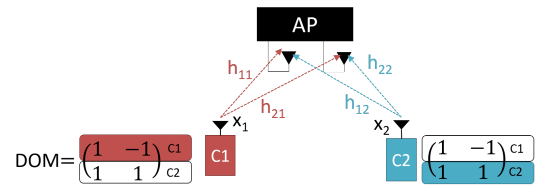

Fig. 4: Channel Estimation with Distributed Transmitters.

The DOM Matrix, as depicted in Fig. 4, permits a single-antenna device to send the required set of training signals that enable CSIR computation at the AP for any MIMO configuration. We use a matrix having rank equal to the total number spatial streams (N_{SS}) with values as specified in the 802.11's orthogonal mapping matrix (see [1] and the example of the figure). However, in contrast to 802.11, each distributed stations select a different row of the matrix to apply as directed by the UO.

C. DSSM MAC

DSSM-MAC enables distributed transmit-processing of uplink transmissions while also disguising the transmission as if it was generated from a single device with a co-located antenna array rather than a distributed array of DSSM nodes.

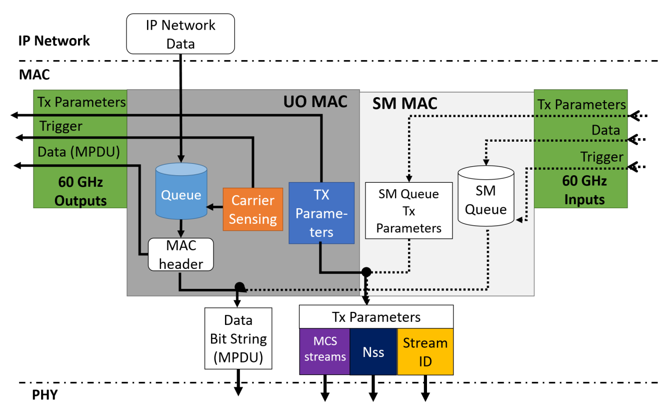

Fig. 5: DSSM-MAC Block Diagram.

Spatial Multiplexing Block. To maintain 802.11 medium access unchanged and to enable medium access isolation among the uplink Originator (UO) client and the Spatial Multiplexers (SM), we design an independent spatial multiplexer block within the device's MAC. The spatial multiplexer block acts as a remote partial-replica of the UO-MAC at the Spatial Multiplexer device. The UO is the only device responsible for carrier sensing for medium access and ensuring MAC and PHY headers emulate a transmission from a single multi-antenna device. Prior to sharing the data to the SMs, the UO adds the MAC header to the payload data, thus sharing the data as an MPDU. This design gives flexibility to the UO to perform MPDU or MSDU aggregation prior to sharing the data to the SMs. To enable PHY header encapsulation on a remote device, the UO forms and shares the transmit parameters defined per stream, which are stored and used remotely by the spatial multiplexer PHY, as shown in Fig. 5. Through this MAC design, we remove PHY dependency, such that the SM-PHY is oblivious of the source of the data and only requires the transmit parameters to process the input data string for transmission. Additionally, to enable flexibility to unknown time that the medium will be obtained, the data and parameters are shared with the SMs in the mm-wave band prior to the medium access attempt.

In DSSM, we grant the UO control of medium access, retransmissions and acknowledgments for the multiplexed data. Consequently, contention is performed *only* by the UO device, not the SM devices. Upon winning contention, the UO triggers the UO queue and all remote SM queues. The trigger releases the bit-string and transmit parameters into the PHY in all transmitting devices. Note that in DSSM having the transmission performed by devices distributed in space increases the medium reservation footprint. Even though common hidden terminal problems that exist in 802.11 random access remain in DSSM (802.11 medium access is unchanged), an increased medium reservation footprint aids hidden terminals.

Multi SM-queue Operation. As illustrated in Fig. 5, the DSSM MAC enables a node to operate as both Uplink Originator and Spatial Multiplexer. The UO-MAC has a primary queue with the device's own traffic and the SM-MAC contains one or more spatial multiplexing queues. This queue separation partitions traffic among OU and SM devices, maintains protocol stack functionality of the DSSM device, and preserves order in which stations obtain medium access.

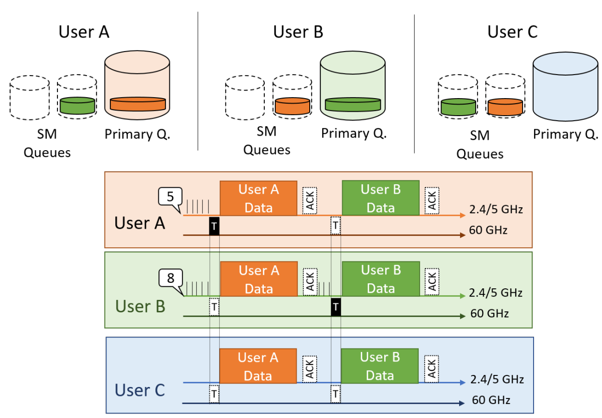

Fig. 6: DSSM-MAC Example.

Fig. 6 depicts an example of the operation of multiple SM-queues with three stations A, B and C, in which two of the stations, A and B, have traffic in their primary queue and all stations enable spatial multiplexing for one another. All stations have previously shared their data through the mm-wave band and each station has two SM-queues.

Once the medium becomes idle, only stations with traffic in the *primary queue* (A and B) select a random backoff counter to compete for medium access. Station A selects a random backoff counter of 5 and station B of 8. Once A's counter expires it triggers the other stations. All three stations transmit the data of station A spoofing an uplink 3x3 SU-MIMO transmission. After the transmission is acknowledged by the AP and the medium becomes idle, B proceeds with its countdown to access the medium. After expiration, B triggers the SM devices and both station A and C join the transmission by transmitting B data. Notably, the order in which stations access the medium is not affected: even if A and B where transmitting via SIMO, they would have transmitted in the same order, thus we maintain the 802.11 medium access unchanged. The key difference is that now more data is transmitted in less amount of time leading to efficient usage of the medium, reduced transmission time and higher throughput.

III. Implementation and Testbed

A. DSSM Implementation

We implement DSSM on a software defined radio platform, WARP, which enables Over The Air (OTA) experiments [2]. We develop evaluation platforms for both legacy and mm-wave bands in hardware as depicted in Fig. 7.

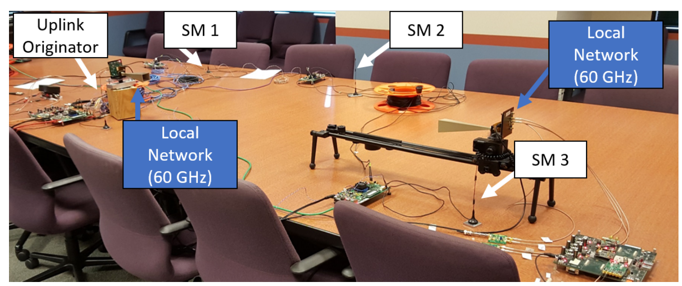

Fig. 7: DSSM Evaluation testbed for legacy and mm-wave bands.

2.4/5 GHz Distributed Transmitter. To achieve standard compatibility in the legacy band, we use the Mathworks WLAN System Toolbox as a building block for DSSM implementation. The WLAN toolbox provides standard-compliant functions for the design, simulation, analysis, and testing of the PHY layer of wireless LAN communications systems [3]. We use the WLAN toolbox as a building block to generate 802.11n compliant signals. We modify and integrate the WLAN toolbox with our implementation of the DSSM PHY in WARPLab for OTA transmission and receiver processing.

Off-the-shelf 802.11n Legacy-Band Receiver. To demonstrate that DSSM's muti-stream transmission is decodable with commodity receivers (spoofing), we perform a subset of experiments with a 2013 MacBook air as the off-the-shelf receiver communicating with two single-antenna DSSM transmitters. This MacBook air uses the Broadcom 802.11n/ac Wi-Fi chip BCM4360 supporting up to 3-streams; however its implementation shows a maximum of 2-streams [4], [5].

DSSM mm-wave Band Testbed. We develop an experimental platform for the mm-wave band consisting of the mm-Wave development platform from Vubiq/Pasternack [6] integrated with WARPLab for the generation of I/Q baseband signals. The mm-wave platform serves as the transmitter and receiver which communicate in 58 to 64 GHz unlicensed channels. To achieve directional multi-Gbps data transmission, we utilize horn antennas with beamwidths of 7, 20 and 80 degrees. Through WARPLab we generate OFDM baseband signals and feed these to the mm-wave platform which upconverts these for OTA transmission to be received by the 60 GHz receiver and down converted to be filtered and demodulated by WARP. Given the 20 MHz bandwidth limitation of the WARPLab baseband signals, we collect signal strength (RMS) traces from the decoded signals and translate these to supported 802.11ad data rates for OFDM PHY based on receiver sensitivity.

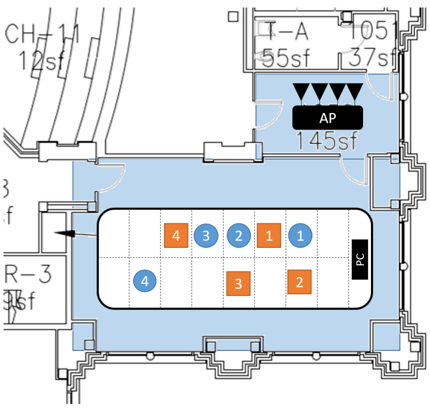

Fig. 8: DSSM experimental setting with two topologies.

Evaluation Scenario and Setup. We perform experiments in a large conference room of 79.6 sq m (857 sq ft) containing two sections that are separated by a wall and door as shown in Fig. 8. For legacy band evaluation, the AP is placed in a small section of the conference room in order to enable NLOS communication with the transmitters placed at the large area of the conference room. We select NLOS communication between transmitters and the AP to represent the scenario in which AP is not reachable with the 60 GHz band.

We evaluate DSSM with two topologies of four transmitters, where for each topology, node placements are randomly selected from a uniform distribution of a 16 square-grid division of the conference table, and transmitters are placed at center of the selected square. The evaluated topologies are shown in Fig. 8, in which the first topology is marked by circle placements and the second by square placements. For all evaluations, the distributed transmitters are synchronized by sharing RF and sampling clocks as well as time synchronization through trigger-cables emulating the OTA triggers.

Finally, we evaluate the performance of the mm-wave band. Due to the limitation of the 60 GHz platform to one transmitter and one receiver, we place the transmit/receive pair at various locations in the conference room representing distances from 0.5 to 3.5 meters.

IV. Evaluation

In this section we evaluate DSSM performance, we first demonstrate our system gains and then study the system limitations.

A. Spoofing an 802.11n Compliant Receiver

DSSM achieves the spatial multiplexing gains of multi-antenna devices by emulating 802.11 uplink multi-stream transmissions with distributed single-antenna devices in a manner that is transparent to the AP. Here, we demonstrate DSSM's ability to spoof an 802.11n off-the-shelf receiver acting as an AP.

The distributed uplink DSSM transmitters consist of two single-antenna WARP radios with shared RF and sampling clocks as well as 0-delay trigger (trigger offsets are separately explored below). We spoof an 802.11n 2x2 (single user) MIMO transmission. The transmission uses a 20 MHz channel with short GI, achieving the maximum supported data rate of 144 Mbps (MCS 15) in a 2-stream configuration. We run wireshark at the receiver to detect DSSM's 802.11n standard compatible transmission, set to monitor mode in the channel of our DSSM transmission.

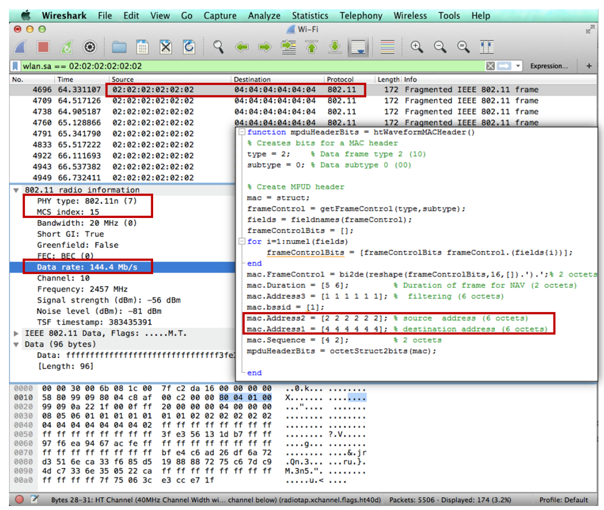

Fig. 9: Demonstration of DSSM spoofing 802.11n spatial multiplexing with off-the-shelf receiver.

Fig. 9 shows a screenshot taken at the receiver showing the received packets from the distributed transmitters in the above configuration. Additionally, Fig. 9 shows the code where the source (02:02:02:02:02:02) and destination (04:04:04:04:04:04) MAC addresses are custom generated and are marked with red on the wireshark screenshot. This confirms that the DSSM transmission is fully standard compatible and that DSSM has "tricked" the 802.11n receiver into believing it was communicating with a single two-antenna device when in fact it was communicating with two distributed DSSM nodes.

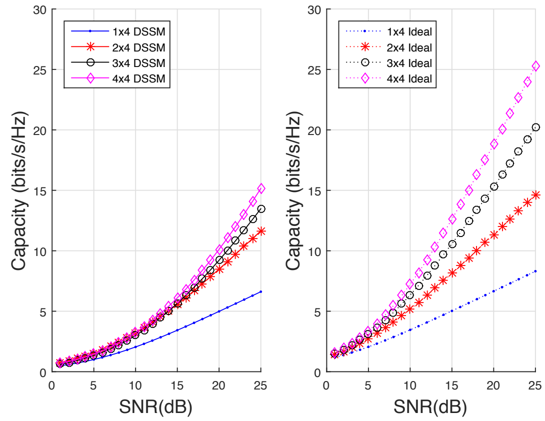

Fig. 10: DSSM Empirical Capacity.

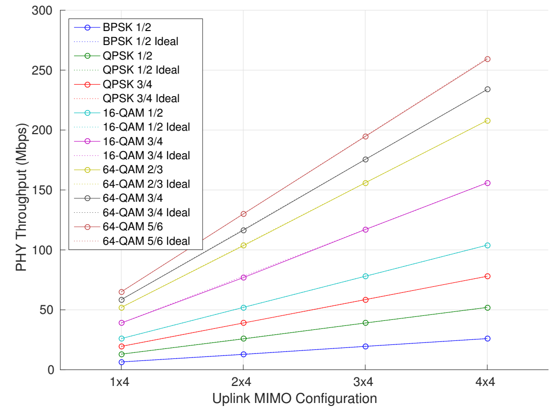

Fig. 11: DSSM Spatial Multiplexing.

B. Scaling and Spatial Multiplexing

Next, we study DSSM spatial multiplexing performance as we increase the number of transmitters. We perform this experiment with our testbed in a large conference room with two topologies, where DSSM devices emulate a 1x4 to 4x4 single-user MIMO uplink transmission to a four-antenna AP. All transmitters have time and phase synchronization in this setup.

Capacity Scaling: To analyze the limits of DSSM, we study the empirical capacity in the evaluated scenario using the generalized Shannon capacity formula for M transmit antennas and N received antennas given by $C(bps/Hz) = \log 2[\det(I_N + (\rho/M)(HH^*))]$ [7], where H^* is the conjugate transpose of the channel H, I_N is the NxN identity matrix and ρ is the average SNR. We use the channels (H) from the OTA experiments, where H is measured at the receiver after implementing DSSM. We calculate the capacity for each subcarrier at each transmission using the measured channels and sweep the average receive SNR values from 0 to 25 dB.

Fig. 10 depicts the average capacity for each MIMO configuration, including both DSSM capacity (left) and capacity for ideal orthogonal channels (right). At high SNR of 25 dB, DSSM achieves 6.37, 11.5, 13.3 and 16.3 bits/sec/Hz for a 1x4 to 4x4 single user MIMO transmission respectively. This demonstrates DSSM achieves up to 2.56x increase in capacity (b/s/Hz) over the SISO case at high SNRs. Fig. 10 demonstrates that even at low SNR, DSSM shows a small increase in capacity. For example, for a value of 5 dB, DSSM achieves 0.88 b/s/Hz for a SISO transmission and 1.70 b/s/Hz for 4x4 spoofed single user MIMO transmission resulting in a 1.93x capacity increase. Based on the empirical capacity, DSSM supports a maximum data rate of 326 Mbps for a 20 MHz channel with an SNR of 25 dB. Next, we observe that DSSM capacity falls below ideal capacity by a factor of 1.5x (e.g., ideal capacity of 25.3 b/s/Hz vs. DSSM capacity of 16.3 b/s/Hz) due to residual channel correlation. Even though DSSM design enables signal decorrelation and achieves large multiplexing gains at high SNRs, residual inter-stream interference may exist due to the OTA-channels not being perfectly orthogonal.

Achievable Data Rates: Next, we evaluate DSSM's ability to achieve 802.11n data rates and demonstrate the spatial multiplexing gains in the evaluated scenarios. Fig. 11 depicts the performance achieved by DSSM for all 802.11n supported modulations when performing a 1x4 to 4x4 MIMO configuration where the transmitter is formed by 1 to 4 single-antenna nodes communicating with a 4-antenna NLOS AP. The x-axis depicts the MIMO configuration (1x4 to 4x4) and the y-axis depicts the achieved performance shown by PHY Throughput which is obtained by subtracting the Bit Error Rate from the expected Data rate for a given MCS.

As targeted, DSSM's performance increases as more spatial multiplexing nodes are used to increment the number of data streams. The multiplexing gains are more pronounced with higher data rates: the highest data rate, 64-QAM 5/6, increases performance from 65 Mbps for 1x4 to 259 Mbps for 4x4. We observe multiplexing gains up to 4x for all evaluated MCS, with a slight decrease where bit errors exist. For example, modulations with 3/4 coding such as 64-QAM 3/4 and 16-QAM 3/4 experience bit errors in the 2x4 case due to interference caused by channel correlation.

C. Robustness to Time Synchronization

Time synchronization among the distributed transmitters plays a key component to emulate a single-node 802.11n multi-stream transmission. A delayed start of transmission from the distributed transmitters leads to Inter-Symbol-Interference (ISI) resulting in issues (explained below) that may cause packet error. A delay in transmission can be caused by several factors such as triggering delay, receive-to-transmit switching delay, and PHY processing delay. In this evaluation we study DSSM susceptibility to time synchronization.

Cyclic Prefix serves as guard intervals typically 400 ns or 800 ns that build tolerance to multipath effects in the channel that spread symbols over time. ISI is introduced when the FFT is taken over a fraction of the desired symbol and the neighboring symbol resulting in decreased effective SNR whose degradation depends on the amount of the neighboring symbol included. In DSSM, delay among transmitters is an additional cause of ISI apart from multipath. Such delay results in EVM increase, eventually leading to stream loss and potential packet loss. As we observe in Fig. 12 transmission delay results in EVM increase eventually leading to stream loss and potential packet loss.

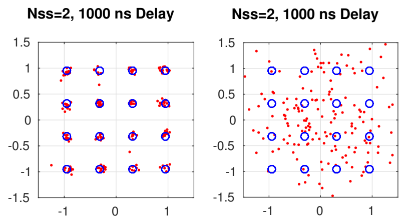

Fig. 12: 2x4 DSSM with 1000 ns delay for second transmitter.

We perform 2x4 uplink 802.11n transmissions with 800 ns guard interval where a UO client and SM client act as a single two antenna transmitter. To isolate the effect of time synchronization, we enable clock synchronization among the transmitters and induce controlled start-of-transmission delays from 0 to 1000 ns to the Spatial Multiplexer stream.

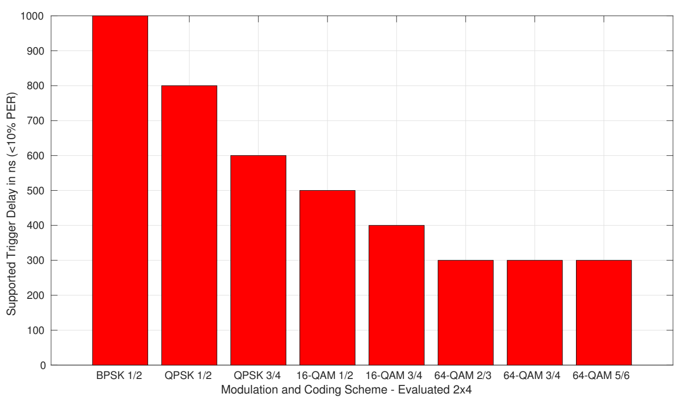

Fig. 13: Robustness to Time Synchronization.

In Fig. 13, the x-axis is the evaluated MCS and the y-axis is the supported delay in nanoseconds for a performance of less or equal to 10% Packet Error Rate (PER). Fig. 13 illustrates three key dependencies of DSSM's performance and MCS that are affected by time synchronization. First, the low MCSs can decode packets even when the delay is larger than the cyclic prefix length (800 ns) even accounting for multipath of OTA transmission. Second, contrary to low MCS, DSSM performance at high MCS degrades before the full cyclic prefix, supporting delays only up to 300 ns for the highest order modulation. Third, while the resilience to trigger delay decreases as the modulation order increases, the degradation stops at a value of 300 ns. The maximum delay value supported by the higher order modulation is dependent on the amount of multipath and delay spread of the evaluated scenario.

The reason behind DSSM's robustness to higher delays than the cyclic prefix length in low order modulation is larger tolerance for EVM; however, the decision regions are smaller for higher order modulation so that the same EVM value results in packet error for higher order modulation but decodable packets for lower order modulation. These experiments demonstrate that precise triggering is required when high order modulation is being used. Likewise, this relationship brings a new factor to MCS selection: the use of peak MCS values is conditioned on synchronization below several hundred nanoseconds.

D. Phase Synchronization

DSSM introduces a new source for carrier frequency offset (CFO), today CFO is a concern between a transmitter and a receiver, however in DSSM the transmission generated from more than one source, thus CFO can also exist between the various transmitters.

In DSSM center carrier frequency of the distributed transmitters may differ as how these drift over-time. Initial CFO estimation is correct due to precise time-synchronization of the transmissions leading to precisely overlapped preambles. However, phase drift of the RF clock from the distributed transmitters may differ and violate the assumption made in the processing of the receiver assuming a single transmitter.

Phase drift across distributed transmitters is studied by downlink multi-user MIMO [8], [9], where these solutions perform transmit beamforming thus have tighter synchronization requirements. Uplink multi-user solutions [10], [11], [12] do not address phase synchronization as no stream or symbol combining is implemented. DSSM does not implement transmit beamforming, however phase synchronization is important as symbols are combined across the transmitters. In this evaluation, we study the robustness of DSSM to phase offset between the transmitters for varying packet lengths.

We perform DSSM 2x4 802.11n standard compatible transmissions with two distributed DSSM clients. We synchronize the clocks of the two transmitters as the desired baseline and induce controlled known phase offset from 0 to 1200 Hz to the second stream. We evaluate various packet lengths, MPDU lengths from 100 to 1000 bytes, for all the modulations supported for two streams. We perform this evaluation with the two topologies presented in §III, where 200 packets are sent for each setting of CFO (7), packet length (11) and Modulation (8).

Fig. 15 depicts DSSM's robustness to phase offset, where the x-axis depicts the phase drift induced to the SM in Hz and the y-axis depicts the supported PSDU length in bytes for a performance less or equal to 10% packet error rate. Fig. 15 demonstrates a the dependency of CFO impact on the tradeoff between packet length and modulation order. A longer packet length is more prone to CFO as the drift over time increases, see Fig. 14. Higher modulations transmit the same amount of bits in a shorter time than lower modulations, thus the same PSDU length results in different packet length based on its modulation. Consequently you expect higher modulations to be more resilient to CFO. However, with higher modulations, symbol decision regions are smaller, resulting in high vulnerability to CFO specially for longer packet lengths.

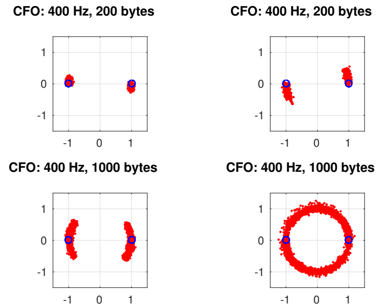

Fig. 14: Effect of 400 Hz CFO in PSDU lengths of 200 (top) and 1000 (bottom) bytes.

This trade off is demonstrated by Fig. 15, where for a given CFO both low modulations and high modulations support lower packet lengths for a threshold of 10% PER, meanwhile modulations such as QPSK 3/2, 16-QAM 1/2 and 16-QAM 3/4 can support longer packet lengths because they represent the balance of this tradeoff. For example for CFO of 400 Hz on the x-axis of Fig. 15, the lowest and highest modulation support only 500 bytes for a 10% PER, meanwhile intermediate modulations (QPSK 3/2, 16-QAM 1/2 and 16-QAM 3/4) can support up to 900 bytes for a 10% PER threshold. Further,

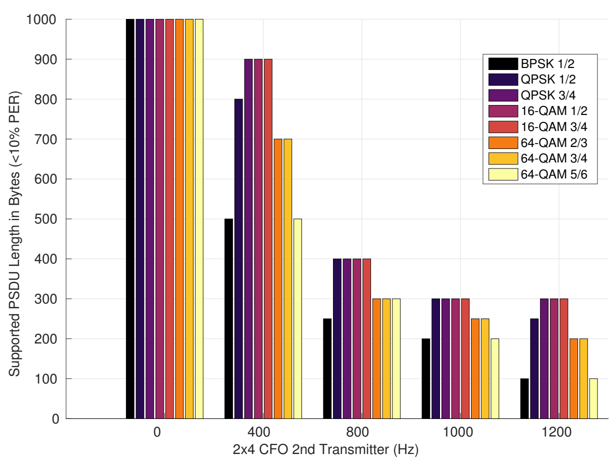

Fig. 15: Impact of Inter-Transmitter's CFO.

Fig. 15 shows DSSM performance rapidly decays as CFO increases, emphasizing the importance of phase synchronization among the transmitters. DSSM provides a control channel among the transmitters for which phase synchronization can be achieved using any of the existing techniques [8], [9]. Phase synchronization can be performed using the 60 GHz control channel if both 2.4 GHz and 60 GHz clocks are referenced on the same oscillator or using the trigger message in the legacy band for in-band synchronization. Through this evaluation we demonstrate the relation DSSM's CFO robustness has to modulation and packet length, where CFO offset introduces a new dimension for MCS selection.

E. SM Stream Outage

mm-wave band link-outage or loss of SM's stream can occur both at legacy or mm-wave bands. Link-outage at the mm-wave band can be caused through blockage, mobility or out-range communication and it is detectable through failure to receive control packets in the mm-wave band. Existing work [13], [14], [15] on mobile device localization aids directional communication with device mobility, these can be integrated to DSSM to increase robustness to SM outage. Stream outage in the legacy can be caused by the node turning off or leaving the network, which can be inferred by outage of the mm-wave link and notified through the mm-wave band if available. Stream outage may be detected prior to legacy transmit opportunity and thus the UO node may be able to adapt. However, in the case SM outage is not detected in time, one or more streams will be lost and may lead to packet failure. In this evaluation we study DSSM robustness to stream outage across varying Modulation and Coding Schemes (MCS) and MIMO configurations.

For this evaluation we perform 2x4, 3x4 and 4x4 SU-MIMO transmissions for all the 802.11n MCSs using a DSSM clients. We evaluated two topologies with outages of one to three streams ensuring one stream (from UO node) is available for communication. Our results showed that a maximum of one stream outage is supported by DSSM to achieve PER of less or equal to 50%. Fig. 16 depicts DSSM robustness to one stream outage. Here, we depict the packet error rate on the y-axis for the three evaluated MIMO settings, for the modulations shown in the x-axis.

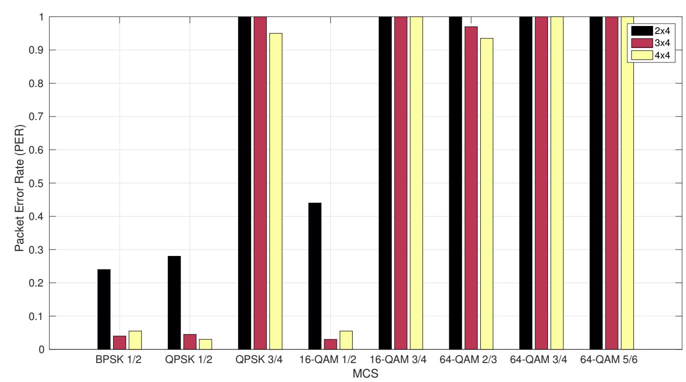

Fig. 16: Packet Error Rate for One Stream Outage.

From Fig. 16 we observe that only 1/2 coding rates are decodable with a stream outage. This illustrates that the decoder is countering the effect of the stream outage, thus only the most robust coding rate of 1/2 can result in a successful packet. Further, we observe that PER is higher for the case of 2x4 MIMO, where even with 50% data loss the decoder is able to recover packets with a PER of 25 to 45%. This exhibits that data loss above 50% is not decodable, explaining why outages of two or three streams are not irreversible by the decoder. That is to say, with one stream outage, DSSM can achieve rates up to 98.29 Mbps for a 4x4, 76.66 Mbps for 3x4 and only 29.12 Mbps for 2x4.

F. Aggregate Airtime

DSSM design targets mm-wave band communication that is an order of magnitude faster than the legacy band. Nonetheless, if the distances or channel qualities between the distributed DSSM nodes are inadequate, the mm-wave band may present a bottleneck. Ideally, with an infinitely fast mm-wave band, the UO communicates with the SM devices instantaneously. However, in the worst case, the 60 GHz links could have sufficiently low throughput that the delay required to share data among the distributed nodes negates any performance benefits of spatial multiplexing, and it would have been preferable to transmit without DSSM using SIMO transmission in the legacy band. Here, we study the effect 60 GHz link quality on DSSM gains characterized by the airtime cost of the packet.

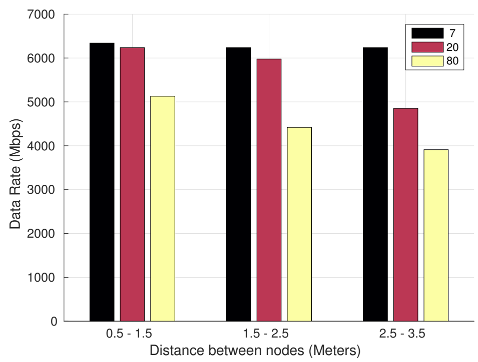

Fig. 17: Experimental mm-wave band data rates.

Experimental mm-wave band data rates. First, we experimentally evaluate the achievable mm-wave band data rates in a conference room for line-of-sight communication between a pair of 60 GHz nodes at 20 different locations for three different receiver beamwidths of 7, 20 and 80 degrees. Fig. 17 depicts the mean data rates in Mbps vs. inter-node distance in meters shown for the three evaluated beamwidths of 7, 20 and 80 degrees. We observe a larger decrease in performance as inter-node distance increases for wider beamwidths, e.g., 80 degree beamwdith performance drops a total of 1218 Mbps. However, for narrow beamwidth such as the evaluated 7 degrees, performance remains close to constant, achieving a minimum of 6.2 Gbps for all evaluated distances. These results demonstrate inter-node communication is an order of magnitude faster than the long-range legacy band and can be maintained over longer distances when narrower beams are used.

Total Air time. To further study the gain of DSSM over 802.11n legacy band SIMO without DSSM, we evaluate the overall airtime cost of sharing data with the spatial multiplexers, or the latency effect of data and parameter sharing. We denote the transmission time for a SIMO transmission as $t_{SIMO} = X/DataRate_L$ where X represents the data length in bits and $DataRate_L$ is the legacy band data rate in bits/sec. For this evaluation we assume all M transmitters use the same legacy data rate.

For DSSM, the overall airtime resources used are defined by:

$$t_{DSSM} = (M-1) * \frac{X}{DataRate_W} + \frac{X/_M}{DataRate_L} \quad (3)$$

where M represents the number of transmitters or streams and $DataRate_W$ is the mm-wave band data rate between the UO and SM nodes. For this evaluation we assume all mm-wave links have equal data rates. Note that the DSSM dual-band transmissions are performed in parallel, such that the mm-wave band transmission time does not have a performance impact on the legacy band, provided that it is less on average (sufficiently small to enable real-time operation).

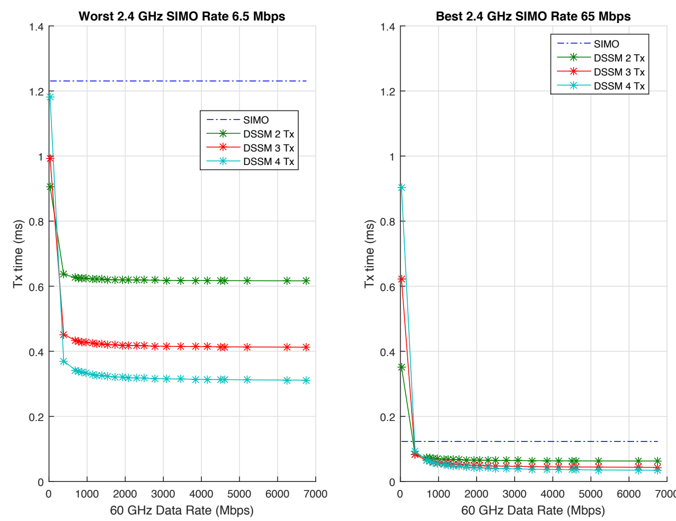

Fig. 18: Total airtime comparison for SIMO vs. DSSM.

Here, we use simulation to compare SIMO and DSSM total airtime with all 802.11ad and 802.11n data rates. The packet length X is 100 bytes and we compare the transmission airtime for DSSM with 2, 3 and 4 transmitters. Fig. 18 depicts the total airtime for DSSM and SIMO vs. the 802.11ad supported data rates. The figure shows that DSSM achieves up to 75% transmit airtime reduction over SIMO even for the highest SIMO modulations. Note that the 60 GHz control PHY base rate (27.5 Mbps) is the only 60 GHz data rate that does not achieve an airtime reduction at high legacy rates such as 65 Mbps. Nonetheless, it does reduce airtime at legacy band base rate (6.5 Mbps). DSSM achieves large airtime reductions because the packet size and thus transmission time in the legacy band is reduced by a factor of M, a reduction that overcomes the latency increase of (M-1) airtime transmissions, in the mm-wave band. Combining the results from both experimental 60 GHz rates and airtime analysis, we can conclude that given that a DSSM link is available, DSSM will improve net airtime, i.e., even when accounting for non-parallel legacy and mm-wave band airtime transmission.

V. Prior Work

DSSM is the first system to enable uplink spatial multiplexing for clients with a single in-band antenna, where the

DSSM client spoofs an unmodified AP to infer that the single-antenna client has an array. Here, we contrast to prior work that enhances *spatial multiplexing* capability in MIMO systems.

Downlink Multi-User WLANs. Antenna asymmetry between the AP and client motivated the standardization of IEEE 802.11ac [16], which transmits to multiple, potentially single-antenna, clients simultaneously in order to obtain a spatial multiplexing gain, typically by zero-forcing among different users and spatial streams [17]. Likewise, while 802.11ac targets a single AP, recent results have shown how spatial multiplexing gains can be scaled among multiple APs [8], [9], [18], [19]. In contrast, DSSM targets *uplink* transmission and obtains gains in a *single-user* paradigm with legacy compatibility. Uplink transmission constraints the design because mobile devices (No. antennas) cannot practically employ multi- or single-user transmit beamforming, distributed uplink transmitters have no infrastructure (wire) connecting them and medium access is decentralized and performed independently by each device in a random access manner.

Uplink Multi-User WLANs. Multi-user transmission has also been proposed on the uplink via successive interference cancellation, interference alignment or orthogonal preambles [10], [11], [12], [20], [21]. In contrast, DSSM employs single-user spatial multiplexing, in which data is generated by a single source and inter-node coordination is performed out-of-band. DSSM removes the overhead for orthogonal user selection [10], [11], [20], [21] and unlike all multi-user schemes, is able to obtain a spatial multiplexing gain even if only a single user is backlogged. Moreover, unlike the aforementioned schemes, DSSM is compatible with the IEEE 802.11n standard.

Cooperative Diversity Cooperative work, e.g. relay systems and Coordinated Multi Point systems, [22], [23], [24], use virtual arrays to increase diversity in low SNR regimes, where the same data is being transmitted through the multiple independent channels (time slots, frequency bands or spatially) to achieve receiver diversity gains. In contrast, DSSM's aim is spatial multiplexing in a high SNR scenario, where each stream carries different symbols of the packet. Unlike relay transmissions with 2 in-band transmissions, DSSM multi-stream transmission is simultaneous (1 in-band transmission). Further, the SMs do not act as relays, instead these process and transmit the data *independently* enabling spatial multiplexing.

VI. Conclusion

In this paper we introduce DSSM, the first system to enable single-user uplink spatial multiplexing for clients with a single-antenna. Without modifying the receiver (an 802.11n compliant AP), DSSM enables a client to spoof a legacy AP to assume the single-antenna client has an antenna array and spatial multiplexing capabilities. We introduce a diverse spectrum architecture that combines high and low frequency bands in a single system to exploit the high data rates of wideband systems to spatially multiplex multiple legacy-band streams. Consequently, we attain both long range communication and spatial multiplexing in antenna-limited devices.

VII. Acknowledgements

This research was supported by Cisco, Intel, the Keck Foundation, and by NSF grants CNS-1642929, CNS-1514285, and CNS-1444056.

References

[1] "IEEE Standard for Information technology Local and metropolitan area networks Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput," *IEEE Std 802.11ac-2014*, Oct 2014.

[2] "WARP project." [Online]. Available: http://warpproject.org

[3] "Mathworks wlan system toolbox." [Online]. Available: http://www.mathworks.com/products/wlan-system/index.html

[4] "The 2013 macbook air review." [Online]. Available: http://www.anandtech.com/show/7085/the-2013-macbook-air-review-13inch/9

[5] "5G WiFi 3-Stream 802.11ac Gigabit Transceiver." [Online]. Available: https://www.broadcom.com/products/wireless-connectivity/wireless-lan/bcm4360

[6] "60 GHz Transmit/Receive (Tx/Rx) Development System." [Online]. Available: http://www.pasternack.com/60-ghz-test-development-system-pem003-kit-p.aspx

[7] G. J. Foschini, "Layered space-time architecture for wireless communication in a fading environment when using multi-element antennas," *Bell Labs Technical Journal*, 1996.

[8] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire, "AirSync: Enabling distributed multiuser MIMO with full spatial multiplexing," *IEEE/ACM Transactions on Networking*, vol. 21, no. 6, pp. 1681–1695, 2013.

[9] H. Rahul, S. Kumar, and D. Katabi, "Megamimo: Scaling wireless capacity with user demand," *Proc. ACM SIGCOMM*, 2012.

[10] K. Tan, H. Liu, J. Fang, W. Wang, J. Zhang, M. Chen, and G. M. Voelker, "SAM: enabling practical spatial multiple access in wireless LAN," in *Proc. ACM MobiCom*, 2009.

[11] A. Zhou, T. Wei, X. Zhang, M. Liu, and Z. Li, "Signpost: Scalable MU-MIMO signaling with zero CSI feedback," in *Proc. ACM MobiHoc*, 2015.

[12] A. B. Flores, S. Quadri, and E. W. Knightly, "A Scalable Multi-User Uplink for Wi-Fi," *Proc. NSDI*, March 2016.

[13] H. Han, S. Yi, Q. Li, G. Shen, Y. Liu, and E. Novak, "Amil: Localizing neighboring mobile devices through a simple gesture," in *Proc. IEEE INFOCOM*, 2016.

[14] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, "Steering with eyes closed: mm-wave beam steering without in-band measurement," in *Proc. IEEE INFOCOM*, 2015.

[15] Y. Ma, X. Hui, and E. C. Kan, "3D real-time indoor localization via broadband nonlinear backscatter in passive devices with centimeter precision," in *Proc. ACM MobiCom*, 2016.

[16] O. Bejarano, E. W. Knightly, and M. Park, "IEEE 802.11ac: from channelization to multi-user MIMO." *IEEE Communications Magazine*, vol. 51, no. 10, pp. 84–90, 2013.

[17] E. Aryafar, N. Anand, T. Salonidis, and E. W. Knightly, "Design and experimental evaluation of multi-user beamforming in wireless LANs," in *Proc. ACM MobiCom*, 2010.

[18] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire, "Achieving High Data Rates in a Distributed MIMO System," in *Proc. ACM MobiCom*, 2012.

[19] Alcatel-Lucent, "Network MIMO." [Online]. Available: http://www.alcatel-lucent.com/wps/DocumentStreamerServlet?LMSG_CABINET=Docs_and_Resource_Ctr&LMSG_CONTENT_FILE=Data_Sheets/Network_MIMO.pdf

[20] K. C.-J. Lin, S. Gollakota, and D. Katabi, "Random access heterogeneous MIMO networks," *Proc. ACM SIGCOMM*, 2011.

[21] T. Wei and X. Zhang, "Random access signaling for network MIMO uplink," in *Proc. IEEE INFOCOM*, 2016.

[22] A. Nosratinia, T. E. Hunter, and A. Hedayat, "Cooperative communication in wireless networks," *IEEE Communications Magazine*, 2004.

[23] A. Sendonaris, E. Erkip, and B. Aazhang, "User cooperation diversity. Part I. System description," *IEEE Transactions on Communications*, vol. 51, no. 11, pp. 1927–1938, 2003.

[24] S. Ma, Y. L. Yang, and H. Sharif, "Distributed mimo technologies in cooperative wireless networks," *IEEE Communications Magazine*, 2011.