

Opportunistic Channel Estimation for Implicit 802.11af MU-MIMO

Ryan E. Guerra
Rice University
ryan@guerra.rock

Narendra Anand
Cisco Systems, Inc.
nareanan@cisco.com

Clayton Shepard
Rice University
cws@rice.edu

Edward W. Knightly
Rice University
knightly@rice.edu

Abstract—Multi-User MIMO (MU-MIMO) linear channel coding can greatly increase wireless system capacity when Stations (STAs) have fewer antennas than the Access Point (AP), but it comes at the cost of significant Channel State Information (CSI) estimation overhead. Previous work has suggested that 802.11af MU-MIMO systems might benefit from long channel coherence time, extending the useful duration of CSI. In this paper, we propose and analyze an opportunistic channel sounding policy that avoids sounding overhead in wireless channels by gathering implicit CSI opportunistically. This policy not only avoids CSI overhead, but also has the potential to enable efficient interoperability of multi-user APs with legacy single-stream STAs. To investigate the performance of this new policy, we implement a new mobile channel sounding framework on a custom 802.11af Software-Defined Radio (SDR) system designed for UHF-band experimentation and evaluate channel sounding performance in indoor and outdoor environments under various mobility modes. Additional protocol analysis shows that in UHF channels with sufficient channel coherence time, an opportunistic channel sounding policy offers significant protocol optimization while improving the scalability of next-generation MU-MIMO systems.

Index Terms—MIMO, MU-MIMO, Beamforming, 802.11af

I. INTRODUCTION

MU-MIMO is a wireless channel coding technique that enables an AP equipped with multiple antennas to transmit *simultaneous* data streams to separate STAs, leveraging spatial diversity to scale data rates with the number of transmit antennas. For an AP to use this technique, it must first estimate the Channel State Information at the Transmitter (CSIT) between each of its transmit antennas and each receiving antenna through a method termed *channel sounding*. The estimated CSIT is then used to compute precoding weights for the multi-stream transmitter. CSIT can also be used for resource allocation, such as user grouping [1] and inter-cell interference mitigation [2].

IEEE 802.11af is a standard amendment for Wi-Fi to operate in unused UHF Television-band White Space (TVWS) channels [3]. The standard can also employ MU-MIMO features of IEEE 802.11ac [4]: here, explicit CSIT is obtained at the AP by first transmitting a sounding packet from the AP to the STA, then having each STA transmit the measured CSI to the AP as a control frame [5]. Unfortunately, the transmission overhead required for CSIT estimation increases with the number of transmit antennas at the AP, M , and the number of aggregate STA antennas, K , and recent results have shown that

this overhead can severely decrease the achievable throughput gains [6], [7].

In this paper, we explore elimination of explicit channel sounding altogether via purely *opportunistic* channel sounding in which CSIT is implicitly estimated from each received *uplink* transmission, whether a data or control frame. Since each uplink frame already contains a training sequence in its preamble (e.g., the TVHTLTF in 802.11af [5]), we use every uplink reception from the STA to the AP, encompassing data, ACKs, and management frames in order for the AP to estimate downlink CSIT. This approach exploits a key property of UHF bands: they can be highly stable on the order of 100 ms while maintaining high multi-user diversity [8]. Thus, the opportunistic policy eliminates CSIT sounding overhead if the channel remains sufficiently unchanged between uplink transmissions.

We show that opportunistic sounding is beneficial in four operating regimes in which: (i) channel conditions are sufficiently stable such that beamforming error due to obtaining CSIT from a prior uplink transmission is negligible; (ii) legacy 802.11 STAs cannot respond to beamforming requests and otherwise could not leverage full spatial diversity; (iii) the number of spatial streams grows such that even implicit channel estimation generates significant overhead; and (iv) the Modulation and Coding Scheme (MCS) is sufficiently high that any wasted airtime due to channel sounding overhead imposes a high relative cost. Scenario (ii) is of particular interest because it enables new 802.11 APs with multi-user capabilities to operate in spectral-efficient multi-user modes with legacy 802.11 equipment that does not otherwise support multi-user modes.

To explore the key performance factors of opportunistic sounding, we design and manufacture a custom MIMO SDR front-end for the WARPv3 SDR platform [9]. This platform enables the first characterization of mobile multi-user UHF channels, enabling evaluation of opportunistic sounding even in the presence of STA or environmental mobility. The design simplifies high-power UHF-band 8x8 MIMO experiments and improves synchronous clocking over previous SDR testbeds. In addition to implementation of custom SDR hardware, we modify a novel SDR channel sounding framework designed for high-speed mobile implicit multi-user channel measurements [10] and port the framework to operate on our UHF equipment.

Finally, we obtain experimental radio licenses WH2XJV and WJ9XFF to operate our experimental equipment on UHF

channels in Houston, TX and perform a series of indoor and outdoor measurement campaigns in various mobility scenarios to analyze MU-MIMO beamforming capacity with respect to CSIT overhead.

We find that fixed wireless nodes utilizing UHF spectrum exhibit long-term stable CSI under environmental and static mobility scenarios. Consequently, we find that with a low number of spatial streams, performance of both active and opportunistic implicit sounding policies significantly exceeds that of the current 802.11af protocol due to the reduced overhead of collecting CSI, even when taking into account the measured beamforming inefficiency of using delayed CSIT. We further extend our analysis to show that opportunistic implicit sounding with more spatial streams yields increasing benefits, enabling future systems with many more antennas than the current maximum of eight in commodity APs.

II. CSIT COLLECTION METHODS

In this section, we present the background necessary for understanding MU-MIMO beamforming, discuss the various approaches to acquiring CSIT and their tradeoffs, and propose a new method for acquiring CSIT opportunistically via implicit channel estimation that will be explored in this paper.

A. Application of CSIT to MU-MIMO

Downlink MU-MIMO is a transmission technique that enables a multi-antenna AP to simultaneously transmit separate data streams to a collection of STAs. This technique is enabled with linear precoding of individual data streams by a collection of complex steering weights. These weights create phase and amplitude-modulated copies of each data stream and simultaneously transmit them from each AP antenna.

Steering weights are represented as a complex weight matrix ($\mathbf{W} \in \mathbb{C}^{K \times M}$), the calculation of which requires knowledge of CSI, the complex magnitude and phase offsets between the transmitter and the receiver antennas represented as the complex channel matrix ($\mathbf{H} \in \mathbb{C}^{M \times K}$), at the transmitter.

A practical method for calculating \mathbf{W} from \mathbf{H} that approaches optimal performance is Zero-forcing Beamforming (ZFBF)[11]. Zero-forcing drives interference between spatial streams to zero, and can be inefficient when users' CSI is not sufficiently orthogonal [12]. ZFBF requires calculation of the \mathbf{H} matrix's pseudo-inverse:

$$\mathbf{W} = \mathbf{H}^\dagger = (\mathbf{H}^H \mathbf{H})^{-1} \mathbf{H}^H, \quad (1)$$

where $(\cdot)^H$ represents the matrix conjugate transpose and $(\cdot)^\dagger$ is the Moore-Penrose pseudo-inverse. When the transmitter precodes with perfect ZFBF weights, \mathbf{W} , signals ideally cancel the effects of the wireless channel at the receiver, allowing each user to receive their own, independent streams.

B. Implicit vs. Explicit Channel Sounding

Implicit beamforming relies on the assumption that the physical channel between the transmitter and receiver is reciprocal in nature so that estimating CSI in the downlink direction is equivalent to estimating CSI in the uplink direction

and vice versa. Accurate array reciprocity calibration has been demonstrated [13], and [14] has experimentally demonstrated equivalent mutual information between downlink and uplink channel estimates utilizing transceiver hardware similar to our own. Therefore, we assume that uplink channel estimation is sufficient to estimate the downlink channel for our purposes, and we assume that all new APs will have the capability to perform reciprocity calibration and provide implicit channel estimation.

The benefits of implicit channel sounding vary based on node/environment mobility as well as the protocol and radio configuration utilized. For example, if the wireless channel varies rapidly due to high mobility, frequent channel sounding, whether implicit or explicit, will be required to obtain accurate CSIT. Per-packet channel sounding mechanisms that incur protocol overhead, such as the multi-user implicit sounding mechanism analyzed in [15] may be required to ensure that channel estimates are accurate in such environments.

However, in the case where the wireless channel remains coherent for long periods of time, for example, due to limited or lack of mobility, then it becomes possible to rely on previously collected CSI for current MU-MIMO transmissions [7], [6]. Practically, such environments exist in wireless networks utilizing sub-GHz carrier frequencies, for instance TVWS networks [8], as well as certain fixed Wi-Fi networks.

C. Opportunistic CSIT Collection

In this section, we propose a new approach to collecting CSIT in 802.11af networks that avoids the overhead of MU-MIMO channel sounding altogether by relying on the opportunistic reception of implicit CSI for regular network traffic.

Fig. 1a diagrams explicit channel sounding, where first the downlink channel is estimated and then the channel estimates are fed back as data packets before each multi-user downlink transmission. This method, with additional polling and channel reservation overhead, is the version currently used in 802.11af [5]. Explicit beamforming overhead scales as $\mathcal{O}(M \cdot K)$.

A proposed implicit sounding method [15] that transmits staggered Null Data Packets (NDPs) in the uplink direction allowing implicit downlink channel estimation following a multi-user trigger from the AP. Its timeline is similar to that of Fig. 1a, but instead the uplink packets are short NDPs rather than Compressed Beam-Forming Report (CBFR) packets, and polling is avoided. Implicit channel estimation overhead is significantly reduced from the explicit case since no CBFR polling or uplink payload is required before the downlink transmission. Implicit beamforming overhead scales as $\mathcal{O}(K)$ since all AP antennas are sounded simultaneously and is key for scaling M , the number of AP antennas.

Fig. 1b displays our proposed *opportunistic* implicit sounding method that estimates the downlink channel implicitly from uplink data transmissions and utilizes that channel estimate for Multi-User Beamforming (MUBF) so long as it remains "fresh." Acknowledgment (ACK) packets from a

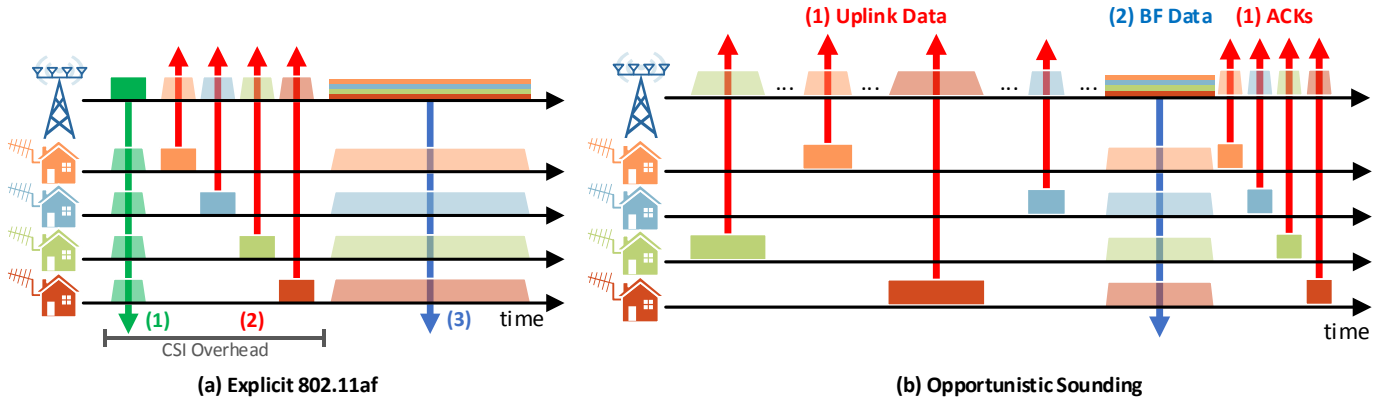


Fig. 1. Time diagram of the considered types of channel sounding. 802.11af polling packets are not included in the explicit sounding diagram.

successful downlink transmission can also refresh CSIT implicitly. Opportunistic implicit beamforming overhead scales as $\mathcal{O}(1)$ since no sounding overhead is injected to sound active STAs.

Given that the UHF channels for 802.11af networks can remain stable for relatively long periods of time, we target to avoid channel sounding altogether and rely on standard PLCP preambles [5] in overheard uplink transmissions to estimate the downlink channel since that estimate will remain valid over multiple packet timescales.

The key strategy for opportunistic sounding is as follows: when historical implicit CSI is available and “fresh,” the AP forms user groups and calculates precoding weights for the optimal multi-user transmission group determined by the MAC scheduler. We utilize two methods when implicit CSI is unavailable or stale for a particular STA: 1) a single downlink frame for the stale STA is de-queued and transmitted by the AP using MISO omni-directional transmission; the subsequent ACK will then provide an update of implicit CSI for that STA; or 2) alternately, the AP could fall back to legacy implicit sounding methods, e.g., [15], if no traffic is available.

In order to determine the feasibility of such an opportunistic sounding policy in 802.11af systems and explore the possible throughput gains, we measure a series of indoor and outdoor multi-user channel traces and perform protocol analysis to understand policy tradeoffs for opportunistic CSIT.

III. EXPERIMENTAL PLATFORM

Although MU-MIMO capabilities will be made available on “wave-two” 802.11ac ASIC chipsets in 2016, the research community is limited by the protocol modifications that can be made to commodity hardware. In addition, no 802.11af ASICs have been announced. For that reason, we have developed the hardware and software stack of a custom SDR platform that allows us to arbitrarily generate, intercept, and modify MU-MIMO transmissions.

Our approach to solving the processing delay in SDR equipment is to take advantage of the higher coherence time of low-frequency channels and perform over-the-air experiments on

selected UHF (470-698 MHz) channels, where the processing latency becomes much less significant [8].

A. Hardware Platform Design

We extend the Wideband UHF Radio Card (WURC) UHF test equipment that we developed in previous work for rapid physical-layer prototyping in UHF bands [8], [16]. WURC was designed to enable high-power transmission up to 1 W and reception of wide-band radio signals in frequencies between 470-698 MHz [16] and each module provides one complete analog radio chain for use with a single WARPv3 SDR base-band [9]. Multiple WARPv3 boards can be clock synchronized with a daisy-chained reference clock and shared sampling trigger.

However, the equipment and daisy-chain topology of [8] suffers from a clocking topology that introduces additional transmission and reception phase errors and aperture jitter as the reference clock signal is forwarded [17]. In addition, it requires the sharing of a transmission/reception trigger over General Purpose Input/Output (GPIO) connectors that also suffers from phase-altering delays and signal bi-stability caused by clock-domain crossing of the trigger signals.

In order to address these issues, we design, layout, and manufacture a clock-synchronized 4-radio adapter board that

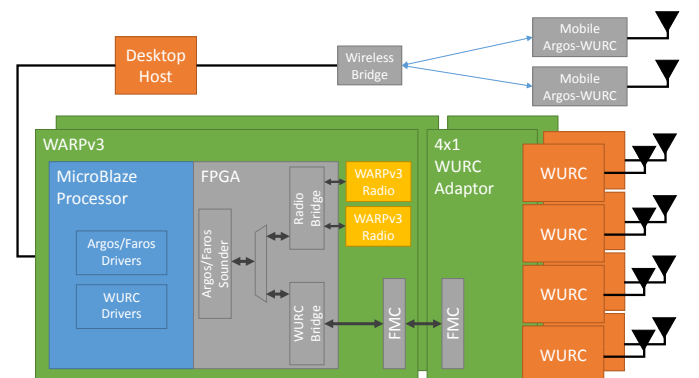


Fig. 2. System diagram of the designed 8×8 Argos-WURC AP and STAs.




Fig. 3. 8x8 UHF-band Argos-WURC array. The 8 red PCBs are single UHF-radio modules, while the green PCBs are the WARPv3 platform and our custom adapter PCB.

can connect up to 4 WURC radio front-ends to a single WARPv3 baseband board (Fig. 2: 4x1 WURC Adapter). With this architecture, baseband sampling and RF reference clocks are now buffered and distributed in a tree topology, while inter-radio triggering is no longer needed since all data streams come from the same FPGA and same clock domain. This has the added advantage of reducing the cost and hardware footprint of a 4-radio UHF AP.

System Configuration. Our over-the-air experiments utilize omni-directional 3 dBi August DTA240 portable UHF antennas for the AP, mobile STAs, and indoor STAs while the static outdoor STAs use Comtelco Y42400WB 7 dBi log-periodic antennas. During experiments, the AP antennas are configured in a linear array separated by a minimum of $\lambda/2$ distance to ensure sufficient spatial diversity.

B. Measurement System: Rapid Implicit Sounding

Channel sounding of mobile devices has always presented a challenge for SDR systems, which require extensive computational resources, power sources, and synchronization to operate as a stand-alone mobile device. For that reason, previous multi-user UHF research focused solely on fixed devices [8] and avoided investigations of channels with nodal mobility.

In order to sound the multi-user environment rapidly between a MU-MIMO AP and multiple mobile STAs, we port the recently published Argos MU-MIMO control channel [10] to inter-operate with our new MIMO WURC array by integrating custom HDL and embedded C libraries for the WARPv3 platform. The robust wireless synchronization scheme utilizing long correlatable signal sequences in Argos allows us to operate mobile STAs remotely, using a low-rate wireless side channel (the wireless bridge in Fig. 2) for STA initialization and control messages, while sample-level synchronization and implicit channel sounding occurs over the UHF channel. More details about the design of Argos are available in [10]; we configure the system to allow us to sound the uplink channel of a set of mobile/static STAs every 2.5 to 5 ms.



Fig. 4. Single-radio mobile node, omnidirectional antenna and outdoor installed AP.

In order to increase the number of antennas available at the AP, we share reference clocks between two 4×4 Argos-WURC APs to create a single 8×8 AP. This significantly shortened clocking topology introduces no measurable decrease in signal or triggering error and has been validated over hours of operation.

IV. EXPERIMENTAL EVALUATION

In this section, we use the data obtained from our indoor and outdoor implicit channel measurements to emulate various implicit resounding policies in several environments.

A. Sounding-Transmission (S-T) Interval

In the following analysis of various alternative resounding policies for 802.11af systems, we focus our analysis on the effect of the time interval between when a channel is sounded and when the final beamformed transmission takes place. We call this time the “Sounding-Transmission Interval,” or S-T interval. Differences in the sounded CSIT compared to the actual physical channel at the time of zero-forcing transmission result in inter-stream interference between STAs as well as reduction in their desired signal strength. In mobile environments, it is highly likely that a larger S-T interval will yield higher inter-stream interference due to increased CSIT error and therefore lower Signal-to-Interference-and-Noise Ratio (SINR).

The S-T interval is important for understanding the performance of opportunistic implicit sounding since an opportunistic AP may have cached, or “stale” CSIT obtained from previous uplink transmissions made at different times. In order to use this CSIT, it will need to make a decision about future beamformed transmissions utilizing that stale CSIT. On the other hand, an implicit or explicit AP refreshes all CSIT simultaneously at the beginning of a multi-user packet, yielding an S-T interval of nearly zero.

Depending on the length of the S-T interval, an opportunistic system could exhibit high inefficiency due to unnecessary sounding overhead, or poor performance due to stale CSIT. In order to emulate opportunistic collection of CSIT, we need to

characterize how drift in the CSI of a single STA will affect the performance of a future beamformed transmission including multiple other STAs.

B. Multi-user Achievable Rate with Increasing S-T Interval

In this section, we investigate the downlink zero-forcing throughput degradation as a function of the S-T interval in indoor and outdoor environments with both nodal and environmental mobility.

Our evaluation methodology relies on the assumption of channel reciprocity. We first record a series of *uplink* channel traces of an 8x4 MU-MIMO system with 4 single-radio STAs using the Argos-WURC system described in §III. This system is used to record multi-user CSI over the course of a minute at regular sampling intervals of 2.5 or 5 ms.

We then assume that the variation in our channel traces is only caused by changes in the physical MIMO channel rather than the radio hardware and use the empirical capacity of the uplink channel in place of the downlink. In [14], the authors demonstrated channel reciprocity using the same transceivers, and in [13] we demonstrated MIMO reciprocity calibration that we have repeated with our hardware from §III-A yet omit here due to space. When accurate reciprocity calibration is performed and interference is identical, the channel capacity in one direction is the same as the other direction.

Each of six different trials was performed either in a: (i) indoor office building environment with non-line-of-sight propagation less than 50 m distance through a wall and a hallway; or (ii) the outdoor heavily forested environment shown in Fig. 4 with non-line-of-sight propagation up to 200 m directly through multiple trees and underbrush. The tested environments were *static*, with no intentional mobility, *environmental motion*, with pedestrians walking around the fixed STAs, or *mobile*, with one (indoor) or two (outdoor) STAs being physically carried by a pedestrian.

ZFBF Rate Calculation. Let P_{jk} represent the signal power of spatial stream j received at STA k . If we let $w_{km} \in \mathbf{W}$ be the transmission precoding weight coefficients from AP antenna m to STA k , and $h_{mk} \in \mathbf{H}$ be the corresponding instantaneous MIMO channel coefficients at the moment of transmission, we can calculate the empirical transmission SINR at STA k as the following:

$$\text{SINR}_k = \frac{P_{kk}}{N_k + \sum_{j,j \neq k} P_{jk}} \quad (2)$$

$$= \frac{|\sum_{m=1}^M h_{km} w_{mk}|^2}{N_k + \sum_{j,j \neq k} |\sum_{m=1}^M h_{km} w_{mj}|^2}. \quad (3)$$

Using the well-known Shannon-Hartley theorem, we calculate the empirical achievable rate of the beamformed channel as $R_k = \log_2(1 + \text{SINR}_k)$.

We compare the loss of achievable per-user throughput in Figs. 5 and 6 as a function of the S-T interval. The zero-forcing achievable rate is the percent difference between the rate with fresh CSIT and the estimated rate using delayed CSIT.

Effect of Mobility on Achievable Rate. We first compare the 8×4 zero-forcing results for the various STAs in Fig. 5.

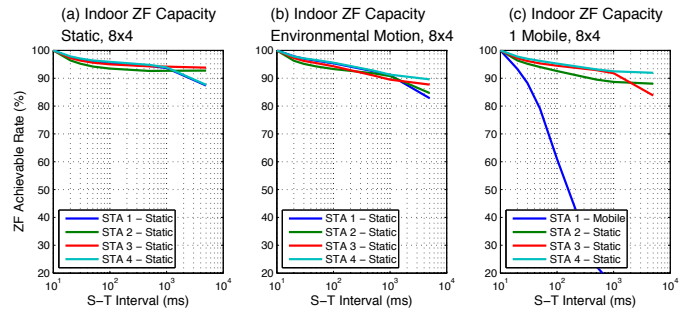


Fig. 5. NLOS Indoor hallway 8x4 UHF-band. Colors correspond to STAs.

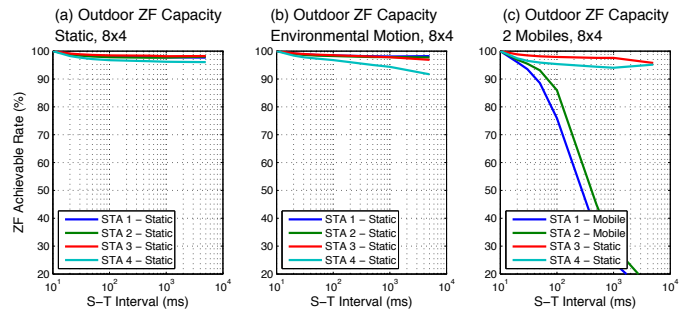


Fig. 6. NLOS Outdoor forest 8x4 UHF-band. Colors correspond to STAs.

When STAs are static, in Fig. 5a and Fig. 5b, we observe that there is minimal loss of beamforming performance as the S-T interval grows. While we would expect that little to no change in CSIT would occur in the largely static environment in Fig. 5a, an unexpected finding is that environmental mobility, even in the non-line-of-sight environment with pedestrians walking within the same hallway, Fig. 5b, had no significant effect on the averaged beamformed rate. Inspecting channel traces, we observe dips in beamforming performance as pedestrians walked by STAs, but such disruptions were small, momentary and had little effect on the average rate, returning to high rate after the pedestrian had passed. Even at 1 second S-T intervals, the system resounds rapidly enough that minimal disruption to the STAs average capacity is observed.

On the other hand, when the STA itself becomes mobile, in Fig. 5c, achievable capacity for the mobile STA dropped quickly after an S-T interval of approximately 20 ms. This still represents a timescale of tens of packets for a mobile STA, indicating that when sufficient uplink traffic is available, an opportunistic sounding AP would provide per-user beamforming performance within 15% of ideal to *mobile* nodes even with S-T interval on the order of a 20 ms. Even under environmental mobility, 15% of ideal beamforming performance would be achieved with a S-T interval on the order of a second.

Effect of Environment on Achievable Rate. We now repeat the same measurements with the same equipment in the outdoor forest environment in Fig. 6. We chose to perform channel sounding experiments in a heavily forested environment since one of the potential applications of 802.11af MU-MIMO networks is to provide last-mile connectivity for

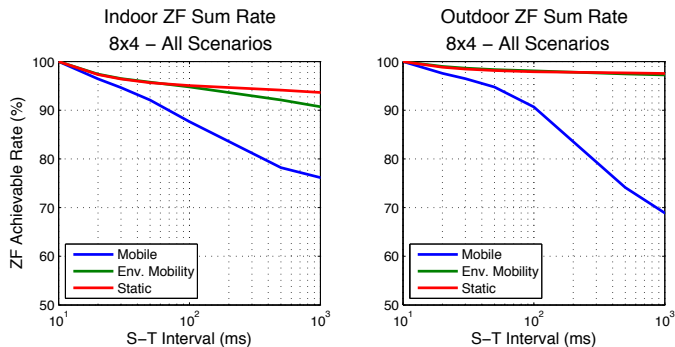


Fig. 7. NLOS 8x4 UHF-band sum-rate.

residential networks in locations where line-of-sight channels are not available for 802.11ac equipment, which also has severe problems propagating through trees [18]. Parallel 2.4 GHz MU-MIMO measurements were attempted at this forested location, yet the signal could barely propagate more than 15 m in the environment and the results were abandoned.

Our results for the forested environment are similar to the indoor environment: as the S-T interval increases, the average supported capacity of the outdoor ZFBF system decreases slowly for the static STAs and much more rapidly for the mobile STAs. A noticeable difference is that the S-T interval breakpoint for the outdoor mobile nodes appears at approximately 50 ms while in the indoor tests it appears around 20 ms. This would be consistent with the outdoor environment that, while also non-line-of-sight, has fewer multi-path reflectors and thus exhibits less channel variation as the STAs move.

We find that based on measured beamforming capacity, up to 1 second of S-T interval is allowable to achieve within 15% of ideal per-user beamforming capacity to fixed STAs, or 20 ms of S-T interval to achieve within 20% of ideal beamforming capacity with mobile STAs in an 8×4 zero-forcing system.

Sum-Rate Results. We sum the individual results obtained in Figs. 5 and 6 in Fig. 7, to report that the sum-rate rate loss with increasing S-T interval is somewhat eased when considering the sum network throughput. This will be used to simulate opportunistic sounding in § IV-C.

Limits of a fixed S-T interval. We evaluate the effectiveness of using a fixed S-T interval to achieve a particular performance level. Vendors of fixed wireless 802.11 equipment are increasingly replacing the 802.11 DCF MAC with Time-Division Multiple Access (TDMA) alternatives for increased long-range efficiency and QoS [19] and could guarantee that opportunistic CSIT is available with a given S-T interval.

Fig. 8 depicts two seconds of the empirical achievable rate of the indoor 8x4 ZFBF system in order to demonstrate the problem of using fixed resounding intervals. The achievable rate of three STAs are shown in different colors; the solid line is the oracle ZFBF rate and the dotted line is the achievable ZFBF rate assuming a fixed 100 ms re-sounding interval. As expected, the mobile STA 1 in blue, which is carried at pedestrian speed within the hallway, demonstrates rapidly changing CSI that cannot be tracked accurately by this large

fixed sounding interval. At each re-sounding point, the periodic system matches the oracle capacity, and then rapidly degrades to approximately 20% of optimal. As the mobile STA 1 physically moves by a static STA 2 (red, 27 seconds), it perturbs its relatively static wireless channel resulting in severe capacity loss.

Such an event is difficult to predict and could result in outages or large capacity loss unless identified and corrected. Based on our observations, a fixed S-T interval would either result in either unnecessary sounding or excessive capacity loss due to stale CSIT since channels can change mobility state rapidly.

Thus, we find that an opportunistic sounding policy should have an adaptive component that adjusts the maximum tolerable S-T interval based on current channel conditions and the mobility state of the STA.




Fig. 8. 8x4 zero-forcing achievable rate, indoor with one mobile STA (blue).

C. Performance of Proposed 802.11af Sounding Alternatives

In this section, we investigate the protocol gains available from an opportunistic CSIT system with regard to various MAC-layer parameters. We also compare performance against an implicit sounding policy adapted from 802.11n standard proposals for multi-user operation [15].

The 802.11af standard attempts to amortize explicit sounding overhead by transmitting aggregated data frames, however the efficiency of this approach depends on the number of frames actually available to aggregate. We analyze the protocol performance of a MU-MIMO system with various channel sounding policies and with varying packet aggregation values in order to emulate both best and worst case scenarios.

We set the single frame size to 1500 bytes, the largest regular Ethernet frame size and the best case for CSIT overhead amortization before aggregation. We compare three different channel sounding policies:

Explicit 802.11af. This is the current standard operation of 802.11af MU-MIMO. CSIT overhead in this case is caused by the NDP Announcement, the sounding NDP, and the sequence of polls and CBFR responses from all 802.11af STAs before each downlink transmission [5]. The upper and lower bounds on explicit performance are calculated with minimum and maximum feedback compression of the CBFR payload, a highly vendor-specific implementation parameter. We assume no impairment on performance from feedback compression,

and plot the median performance while indicating the bounds with a shaded red region.

Although the 802.11af standard only supports up to 8 concurrent spatial streams, we assume that timing and protocol performance scales with the number of streams in order to provide a point of reference for scaling to large numbers of antennas. We label this policy “*Explicit 802.11af*” in the following plots.

Implicit Proposal for 802.11af. In [15], the authors proposed an alternative multi-user CSI sounding protocol that avoids the lengthy CBFR by estimating the channel implicitly with short NDPs. CSIT overhead in this case comes from the NDP Announcement and a staggered sequence of uplink NDPs that are used for implicit channel estimation before each multi-user transmission as proposed in [15]. Since the channel is estimated implicitly, there are no levels of feedback compression to display. We label this policy “*Implicit*” in the following plots.

Opportunistic Proposal for 802.11af. In this case, there is no CSIT overhead to multi-user transmissions. We explore three regions of operation for an opportunistic AP:

- 1) “*Opportunistic.*” The best-case performance assuming all CSIT is available opportunistically and there is no beamforming penalty for using stale CSIT.
- 2) “*Opportunistic with Bootstrap.*” An alternative fallback mode where at most one STA has stale CSIT and the AP sends a single packet to that STA before each multi-user transmission in order to implicitly refresh its CSIT. This can be viewed as a way of quickly bootstrapping opportunistic CSIT to a STA that previously was inactive.
- 3) “*Opportunistic with Stale CSIT.*” A trace-driven lower bound on opportunistic performance based on our environmental measurement traces. We assume that CSIT is refreshed opportunistically every second. According to our empirical results in Fig. 7, this would result in less than 10% reduction in achievable sum-rate in an environment with static STAs. Thus, we reduce the throughput of the best-case opportunistic scenario by the requisite amount, presenting a more fair approximation of how an implemented opportunistic system might perform.

All ACKs are staggered as per the 802.11af specification. For tractability, transmissions are assumed to be successful, requiring no retransmissions, and only downlink data flows are considered.

1) *Sounding Policy Performance: 4x4:* In Fig. 9, we vary the multi-user frame aggregation number from 1 to 64 for the lowest (top) and highest (bottom) 802.11af MCS in a 4×4 system where all STAs have only a single antenna.

Effect of Frame Aggregation. Frame aggregation allows the cost of channel sounding to be amortized over large payloads. While we expect that increased aggregation will generally decrease the efficiency of channel sounding reduction protocols, it also determines crossover points in terms of protocol performance.

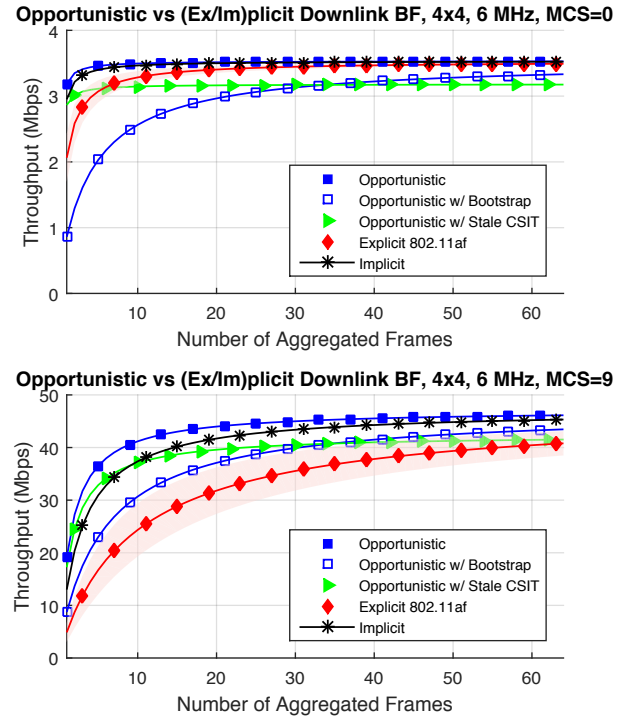


Fig. 9. Network throughput for a 4x4 802.11af system on a 6 MHz UHF channel. Base rate (top) and maximum MCS (bottom).

At the lowest MCS in the top plot of Fig. 9 with no frame aggregation, there is a moderate performance gap between implicit channel sounding methods (opportunistic, implicit) and the current explicit 802.11af policy. An opportunistic sounding policy would increase throughput at best by 31%, while an implicit sounding policy would increase throughput by 21% over explicit 802.11af. However, as the aggregation rate increases, these alternatives rapidly converge.

Effect of MCS. In essence, the lower the MCS rate, the lower the relative overhead of sounding; thus the sounding mechanism matters much less at low rates than high rates.

At base rate, shown in Fig. 9 (top), there is little advantage in opportunistic sounding, and our proposed bootstrapping method under-performs even explicit sounding. However, as the MCS of the system increases, the relative cost of sounding overhead also increases since airtime becomes more valuable, potentially compensating for the stale CSIT penalty.

In Fig. 9 (bottom), we show the same results for the maximum supported MCS. 802.11af sounding overhead is much more costly when the system could otherwise be operating at high MCS, since CBFRs, polling packets, and ACKs are all sent at base rate for robustness. The large range in explicit 802.11af sounding performance (red region) stems from the fact that uncompressed CBFR packets take a significant amount of airtime, resulting in very high overhead. At high MCS, opportunistic sounding can improve throughput by 186% and implicit sounding can improve by 94% without frame aggregation.

While opportunistic sounding with stale CSIT is strictly better than explicit 802.11af up to 35 aggregated frames, it

barely out-performs implicit sounding at low aggregation with fewer than 10 frames and then performs significantly worse with higher frame aggregation.

Therefore, we conclude that for a low number of spatial streams, opportunistic channel sounding has approximately equivalent performance compared to implicit channel sounding and potentially worse performance when considering beam-forming error from stale CSIT. However, both opportunistic and implicit channel sounding offer significant throughput gains over the current explicit 802.11af standard.

The best usage scenario for opportunistic sounding in this regime is when implicit STA cooperation is not possible, such as with current 802.11 devices. A system design that leverages this observation would utilize *opportunistic* CSIT when per-user downlink traffic queues are below 3-52 MB, depending on the current MCS, and then revert to explicit sounding when queues exceed that size and sounding overhead can be sufficiently amortized. For legacy 802.11a/b/g/n devices that do not report any CSIT, only opportunistic CSIT would be available and the decision is made between multi-user and single-user transmission modes only.

2) *Scaling to 32x16*: At all MCS, the challenge of efficiently using narrow bands of UHF radio spectrum is clear: system throughput is no more than 50 Mbps even with full 4x4 spatial diversity at the maximum MCS (Fig. 9). For this reason, we explore the possibility of leveraging additional spatial streams for UHF-band communications as a means of increasing spectral efficiency.

Given the potential for large-scale 802.11af system installations to establish long-range point-to-multi-point networks and the need to support high throughput over narrow UHF channels, we extend our beamforming protocol analysis to a 32×16 system in Fig. 10. Previous work on many-antenna MU-MIMO systems has proposed implicit channel sounding as a means to avoid protocol collapse as the number of antennas at the AP grows [13]. Our results in § IV-B indicate that the CSI of stationary STAs in both indoor and outdoor environments remain constant for long periods of time, which supports the possibility of using opportunistic sounding policies to increase system throughput even further.

In all cases with a large number of spatial streams, explicit channel sounding suffers severely from protocol congestion due to the high number of spatial streams and amount of explicit data that is transmitted to the AP to report CSI.

In Fig. 10 (top), we see that for low MCS rates and frame aggregation below 18 frames, opportunistic sounding with stale CSIT out-performs even implicit sounding, given the number of STAs involved in each transmission.

In Fig. 10 (bottom) at the maximum supported MCS, strict relationships emerge between the sounding policies, since CSIT overhead dominates any other effects at this scale. When channel sounding becomes extremely expensive, the use of opportunistic CSIT is able to offer significant throughput gains over implicit sounding, ranging from 112% with no frame aggregation, to 18% at maximum aggregation, even when

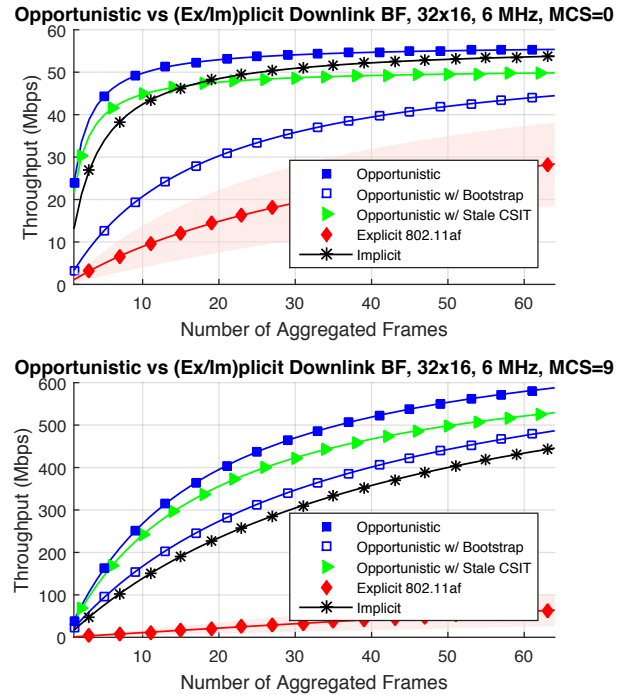


Fig. 10. Network throughput for a 32x16 802.11af system on a 6 MHz UHF channel.

considering the penalty from stale CSIT. Explicit sounding should be avoided altogether.

V. RELATED WORK

Managing protocol overhead is crucial for achieving multiplexing gains with MU-MIMO transmissions and thus numerous works propose techniques that seek to lessen the impact of obtaining CSIT.

Overhead amortization. One category of work seeks to reduce CSI overhead through the use of techniques such as frame aggregation that amortize overhead across multiple data frames. For example, the authors of [20] develop custom frame aggregation techniques for amortizing explicit overhead in 802.11ac systems. Additionally, works such as [21] concede that the overhead of CSIT acquisition is so detrimental, that they suggest avoiding MU-MIMO transmissions altogether.

Our protocol simulations considering both aggregation and compression in Fig. 9 and 10 demonstrate that implicit sounding offers significant benefits while opportunistic sounding can improve even further in certain cases.

Channel sounding suppression. The second category of techniques seek to reduce CSI overhead by avoiding channel sounding when possible. For example, MUTE [7] reduces explicit sounding overhead by opportunistically sounding users when the wireless channel is free and by tracking channel variation to avoid sounding the channel unnecessarily. In our work, we focus on more stable 802.11af channels and higher-order MIMO systems where it becomes feasible to avoid sounding altogether and maintain the same throughput performance as with full explicit channel sounding, while scaling well. Our approach further enables ZFBF to STAs without 802.11ac/af CSI reporting enhancements.

AFC [6] proposes a protocol that allows STAs to determine their own downlink CSI variation through the use of a “Compression Noise” (CNo) metric which tracks the difference in CSI measurements over time and only requests sounding when needed. We propose and analyze an alternative opportunistic approach that avoids explicit sounding altogether.

Implicit channel sounding. Precoding schemes rely on CSIT provided by the radio physical layer. Previous work shows that under many conditions, the additional explicit protocol overhead [22], [7] and CSI feedback compression error [15], [6] in explicit channel sounding can severely degrade the performance of the 802.11ac MU-MIMO protocol. A case where implicit CSI is not only beneficial but necessary is when the number of antenna on a given wireless device grows large, such as in “massive” or many-antenna MIMO, where explicit sounding cannot scale efficiently [13].

While not the first to propose implicit channel sounding, we are the first to measure the beamforming and protocol cost associated with various channel sounding techniques and to propose a completely sounding-free approach for fixed wireless systems with long coherence time.

CSIT Prediction. Other work has focused on attempting to predict CSIT from historical measurements for adaptive modulation systems [23]. When deciding when to use opportunistic CSIT or when to fallback to other sounding modes, knowledge of the expected cost of stale CSIT is crucial. Our results in Fig. 6 demonstrate that beamforming degradation can vary from STA to STA based on their mobility state, therefore future work might focus on extending CSIT prediction algorithms for beamforming and sounding mode selection.

VI. CONCLUSION

In order to scale the capacity of MU-MIMO beamforming systems, it is important to address the problem of CSIT overhead, particularly in 802.11af systems with potentially limited system bandwidth.

In this work, we developed a new SDR system specifically for UHF-band MU-MIMO that allowed us to gather the first mobile multi-user channel traces in the UHF band, which can be found in [24]. Based on our analysis of beamforming capacity with stale CSIT, we showed large S-T intervals can be tolerated in UHF frequency bands, enabling the gathering of CSIT purely opportunistically and enabling multi-user transmissions with legacy 802.11 equipment that can not provide CBFR reports.

We compared three different channel sounding policies and showed that for a small number of spatial streams, significant throughput gains are available with either of the implicit sounding policies, though the penalty of using stale CSIT would encourage the use of implicit sounding rather than opportunistic sounding, if available. However, as the number of spatial streams increases, the overhead of even implicit beamforming begins to become a bottleneck on 802.11af performance and opportunistic channel sounding becomes much more beneficial.

ACKNOWLEDGEMENTS

This research was supported by Cisco and by NSF grants CNS-1514285, CNS-1444056 and CNS-1126478.

REFERENCES

- [1] J. Mao, J. Gao, Y. Liu, and G. Xie, “Simplified semi-orthogonal user selection for MU-MIMO systems with ZFBF,” *Wireless Communications Letters, IEEE*, vol. 1, no. 1, pp. 42–45, 2012.
- [2] M. Rahman and H. Yanikomeroglu, “Enhancing cell-edge performance: a downlink dynamic interference avoidance scheme with inter-cell coordination,” *Wireless Communications, IEEE Transactions on*, vol. 9, no. 4, pp. 1414–1425, 2010.
- [3] A. B. Flores, R. E. Guerra, E. W. Knightly, P. Ecclesine, and S. Pandey, “IEEE 802.11 af: a standard for TV white space spectrum sharing,” *Communications Magazine, IEEE*, vol. 51, no. 10, pp. 92–100, 2013.
- [4] IEEE Std 802.11ac-2013, “Amendment 4: Enhancements for very high throughput for operation in bands below 6 ghz,” 2013.
- [5] IEEE Std 802.11af-2013, “Amendment 5: Television white spaces (twvs) operation,” 2013.
- [6] X. Xie, X. Zhang, and K. Sundaresan, “Adaptive feedback compression for MIMO networks,” in *Proc. ACM MobiCom*, pp. 477–488, 2013.
- [7] O. Bejarano, E. Magistretti, O. Gurewitz, and E. Knightly, “MUTE: Sounding Inhibition for MU-MIMO WLANs,” *Proc. ACM SECON*, 2014.
- [8] N. Anand, R. E. Guerra, and E. W. Knightly, “The Case for UHF-band MU-MIMO,” in *Proc. ACM Intl. Conf. on Mobile Computing and Networking (MobiCom)*, (Maui, HI), Sept. 2014.
- [9] “Warp project.” <http://warpproject.org>.
- [10] C. Shepard, A. Javed, and L. Zhong, “Control Channel Design for Many-Antenna MU-MIMO,” in *Proc. ACM MobiCom*, MobiCom ’15, (New York, NY, USA), pp. 578–591, ACM, 2015.
- [11] T. Yoo and A. Goldsmith, “On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming,” *IEEE Journal on Selected Areas in Communications*, vol. 24, Mar. 2006.
- [12] 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*, (Chicago, IL), Sept. 2010.
- [13] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong, “Argos: Practical many-antenna base stations,” in *Proc. ACM MobiCom*, (Istanbul, Turkey), Aug. 2012.
- [14] M. Guillaud and F. Kaltenberger, “Towards practical channel reciprocity exploitation: Relative calibration in the presence of frequency offset,” in *Proc. IEEE WCNC*, pp. 2525–2530, 2013.
- [15] H. Lou, M. Ghosh, P. Xia, and R. Olesen, “A comparison of implicit and explicit channel feedback methods for MU-MIMO WLAN systems,” in *IEEE PIMRC*, pp. 419–424, 2013.
- [16] “WURC Wideband UHF Radio Card.” <http://www.skylarkwireless.com/WURC>.
- [17] W. Kester, “Aperture time, aperture jitter, aperture delay timeremoving the confusion,” *Analog Devices, MT-007 Tutorial*, Oct. 2008.
- [18] G. Durgin, T. S. Rappaport, and H. Xu, “Measurements and models for radio path loss and penetration loss in and around homes and trees at 5.85 ghz,” *IEEE Trans. on Comm.*, vol. 46, no. 11, pp. 1484–1496, 1998.
- [19] S. Unni, D. Raj, K. Sasidhar, and S. Rao, “Performance measurement and analysis of long range wi-fi network for over-the-sea communication,” in *WiOpt 2015*, pp. 36–41, IEEE, 2015.
- [20] B. Bellalta, J. Barcelo, D. Staehle, A. Vinel, and M. Oliver, “On the Performance of Packet Aggregation in IEEE 802.11ac MU-MIMO WLANs,” *IEEE Communications Letters*, Jul. 2012.
- [21] A. Thapa and S. Shin, “A MAC protocol to select optimal transmission mode in very high throughput WLAN: MU-MIMO vs. multiple SU-MIMO,” in *Proc. IEEE AH-ICI*, Nov. 2012.
- [22] N. Anand, J. Lee, S.-J. Lee, and E. W. Knightly, “Mode and User Selection for Multi-User MIMO WLANs without CSI,” in *Proc. IEEE INFOCOM*, IEEE, 2015.
- [23] A. Duel-Hallen, “Fading channel prediction for mobile radio adaptive transmission systems,” *Proceedings of the IEEE*, vol. 95, no. 12, pp. 2299–2313, 2007.
- [24] C. Shepard, A. Javed, R. Guerra, J. Ding, and L. Zhong, “Many-Antenna MU-MIMO Channel Measurements,” *Submitted to IEEE Asilomar Conference on Signals, Systems, and Computers*, 2016.