Understanding Real Many-Antenna MU-MIMO Channels

Clayton Shepard[†‡], Jian Ding[†], Ryan E. Guerra[†‡], and Lin Zhong[†]
[†] Department of Electrical and Computer Engineering
Rice University, Houston, TX
[‡]Skylark Wireless LLC, Houston, TX
cws@rice.edu, jian.ding@rice.edu, ryan@skylarkwireless.com, lzhong@rice.edu

***Abstract*—We conducted one of the most comprehensive many-antenna MU-MIMO channel measurement campaigns ever reported. Our measurement system supports full mobility with very high time-frequency resolution. We report channel traces in the UHF, 2.4 GHz, and 5 GHz bands, in diverse environments, with up to 104 base-station antennas serving 8 users. We characterize channel stability and capacity across these bands and environments with varying mobility. Our results show that channels exhibit large capacity fluctuations on the order of milliseconds, even with just pedestrian mobility. However in stationary environments, MU-MIMO channels are long-term stable regardless of frequency, and are only ephemerally affected by environmental mobility. Furthermore, these traces and measurement system are made openly available online, enabling future research on many-antenna wireless systems.**

I. Introduction

Multi-User MIMO (MU-MIMO) has been recently adopted by commercial 802.11 WiFi access points and LTE base stations, providing potential gains of up to 4-fold over their Single-Input, Single-Output (SISO) counterparts by exploiting spatial multiplexing. It has been experimentally demonstrated in controlled propagation conditions [1]–[3] that by adding more antennas to the base stations and increasing the number of simultaneous users, MU-MIMO gains can be further increased by 12-fold or more. However, these gains are highly dependent on the MU-MIMO channels, which are significantly affected by user mobility, propagation environment, and frequency. Such MU-MIMO channels, especially those with a large number of base-station antennas, are not well understood experimentally.

To better understand these MU-MIMO channels in the real-world, we implemented a realtime wideband many-antenna MU-MIMO channel measurement system. Built on the ArgosV2 [4] platform and Faros control channel design [5], this system enables very reliable high time-frequency resolution measurements, supporting sub-millisecond sounding intervals with 20 MHz bandwidth, across the UHF, 2.4 GHz, and 5 GHz bands. We leveraged this system to conduct one of the most extensive and diverse mobile MU-MIMO measurements campaign ever reported, already containing over 1 billion channel measurements on more than 20 topologies, and continuing to expand. These topologies include line-of-sight (LOS) and non-line-of-sight (NLOS) scenarios in both indoor and outdoor environments with various degrees of mobility and multipath. Additionally we constructed an open-source Python channel analysis toolbox to study the fundamental properties of many-antenna MU-MIMO channels.

Our analysis of these measurements provides real-world reference points and reveals important trends that have a significant impact on MU-MIMO system design. While many of these results are intuitive, this data and analysis provide a ground-truth across a diverse number of environments, mobilities, and frequencies. In this paper we focus on a few key points. (*i*) Even under pedestrian mobility, channel stability is very low, e.g., we measured channel coherence times of 16 ms and 7 ms for 2.4 GHz and 5 GHz, respectively, in an indoor NLOS environment. This has important implications for the system design in next-generation wireless systems, as collecting Channel State Information (CSI) every 10 ms can result in a loss of over 50% of the achievable capacity. (*ii*) Multipath, frequency, number of base-station antennas, and rotational mobility significantly affect channel stability, causing up to orders of magnitude change in channel coherence. This suggests that next-generation MU-MIMO systems will need much finer-grained customization than current MU-MIMO systems to operate efficiently under diverse environments, topologies, and levels of mobility. (*iii*) We found that stationary users have indefinitely stable channels, with mobility caused by pedestrians in the environment itself having minor, and ephemeral, impact. Therefore in fixed topologies MU-MIMO systems can virtually eliminate channel sounding overhead. (*iv*) Our measurements also demonstrate that in realistic scenarios channel stability is *bimodal*, that is, users are either mobile or stationary, and thus either have very unstable or very stable channels, respectively. This presents a significant opportunity to reduce system overhead with dynamic channel sounding protocols. (*v*) Finally, we found that channel capacity can fluctuate an order of magnitude faster than the channel coherence. Thus channel coherence, as measured by expected correlation, is *not* a good indicator of the channel resounding interval, though it is often used synonymously, e.g., [1], [6], [7]. Next-generation many-antenna MU-MIMO systems must carefully take into account the characteristics of real-world channels, particularly mobility, to achieve their potential capacity gains.

We have released the measurement system, channel analysis toolbox, and channel measurements online [8], with the hope that they will help guide next generation MU-MIMO system design and analysis.

II. Background and Related Work

Massive MIMO is a key candidate technology for 5G and other next-generation wireless technologies [9]. While it has already demonstrated the potential to increase wireless throughput by an order of magnitude in real-world experiments [1]–[3], it still faces key implementation challenges. In particular, since MU-MIMO relies on forming directional beams towards users, mobility can significantly reduce system performance, which was analyzed in our previous work [10]. Indeed, if not carefully designed, the capacity of a MU-MIMO system can be reduced to less than a traditional SISO system, due to the overhead of realtime CSI acquisition and computation. Until now, the impact of mobility on many-antenna MU-MIMO systems has not been thoroughly characterized in the real world.

To the best of our knowledge, we are the first to characterize high-resolution (as fast as sub-millisecond collection interval) temporal behavior of many-antenna (up to 104×8) MU-MIMO channels, across multiple frequency bands. A number of prior MU-MIMO channel measurements have been reported, including [1], [11]–[23]. The evaluation of empirical time-variant channels has mostly focused on small-scale systems with up to 16 base-station antennas [11]–[18]. Some results from previous work differ from ours, e.g., [16], [17], reported stationary users having unstable channels, whereas our measurements indicate stationary users typically have very stable channels, regardless of environmental mobility.

Many-antenna MU-MIMO channel measurements have been reported in [19]–[23] with up to 128 base-station antennas. In particular, the authors in [23] measured channels with user mobility on a system with 128 base-station antennas and 8 single-antenna users. These measurement campaigns built a comprehensive foundation for realistic many-antenna MU-MIMO channels in static environments, however they do not report analysis or results regarding environmental or user mobility. Since these measurement platforms leverage virtual antenna arrays or multiplexed antenna arrays, which have inherent RF switching overhead, it seems they are unable to achieve the time-resolution required to accurately sample mobile channels. Our measurements indicate that channels for all users should be collected within 100s of μs, and the resounding interval needs to be on the order of a few ms to provide complete and accurate measurements for NLOS environments with pedestrian mobility.

III. System Design

We designed and implemented a realtime wideband many-antenna MU-MIMO channel measurement system that supports high time-frequency resolution across the UHF, 2.4 GHz, and 5 GHz bands. We built this system on the ArgosV2 platform [4], based on WARPv3 [24], and leveraged the Faros control channel design [5] to provide time-frequency synchronization with the users and collect CSI. To support UHF, we ported Argos and Faros to the WARP-based WURC platform [16], which involved converting and merging code, standardizing control interfaces, as well as implementing an

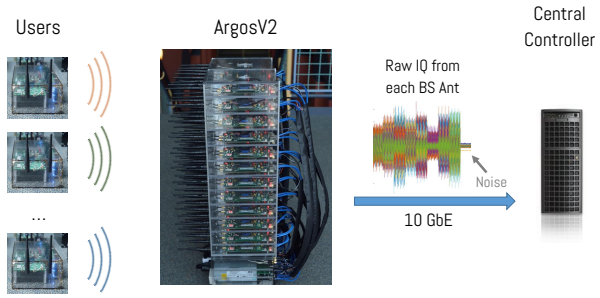

Fig. 1. Overview of the channel measurement system design. At the beginning of each frame the base station sends a beacon to synchronize the users. Each user then sends orthogonal pilots, which each base-station WARP node records. At the end of the pilot phase, base-station nodes report the raw pilot IQ samples to the Central Controller, which records them to an HDF5 file.

IQ only Automatic Gain Control (AGC), as the WURC radio does not have an Received Signal Strength Indicator (RSSI) indicator. Using WURC-enabled battery-powered ArgosMobiles as users, this measurement system can easily be configured to collect fully-mobile channel traces in the UHF, 2.4 GHz, and 5 GHz bands simultaneously, though each band requires separate radios and antennas.

To enable realtime channel measurements, we implemented a from-scratch Python framework that provides complete control over the Argos radio modules, including parameters such as frequency, number of base-station antennas, number of users, frame length, pilot length, etc. This framework is very similar to the WARPLab [24] framework, which is a Matlab interface that provides arbitrary control of WARP nodes over Ethernet. However, our framework is much faster and utilizes asynchronous I/O to simultaneously interface with all of the nodes, as well as handle node-initiated communication. Notably, our Python framework is actually compatible with the default WARPLab design, and they can even be used together in the same experiment.

Leveraging this framework, we implement the Argos CentralController, as defined in [1], and provide functions such as uplink AGC and centralized beamforming computation. Additionally, the CentralController saves raw uplink pilot IQ samples in realtime to an HDF5 file [25], [26], enabling longitudinal traces lasting for hours, or even days, only limited by storage capacity. Optionally, the base-station nodes can also be configured to report uplink data IQ samples, which can also be stored to the same HDF5 file. This trace storage is performed either by a completely native Python implementation, or an accelerated C shared library, which supports sub-millisecond frame lengths, and therefore CSI resounding interval. Due to channel reciprocity, the uplink channels collected are equivalent to downlink channels, only varying by a constant relative phase shift, transmit power, and any asymmetric noise [1].

The current system supports up to 20 MHz bandwidth and 16 simultaneous users. The number of users supported can easily be expanded, at the expense of time or frequency resolution. The number of base-station antennas supported is virtually unlimited, and is reliant only on the computational

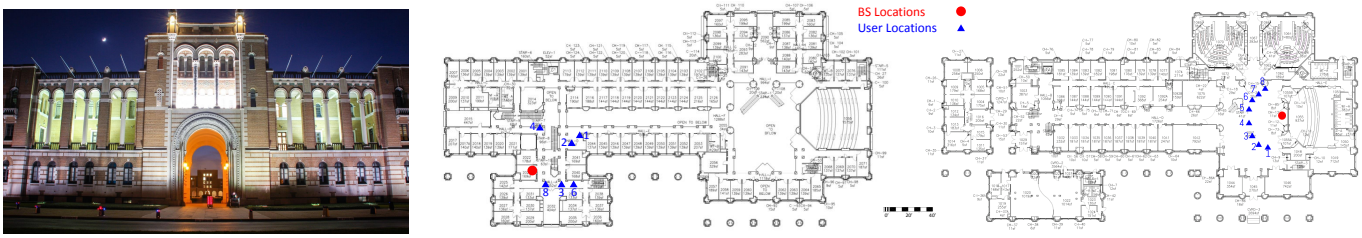

Fig. 2. Example propagation environments. Maps and pictures of each experimental setup are included online with each trace. *(Left)* Picture of outdoor environment. *(Center)* Map of example indoor NLOS environment (users not shown are on other floors). *(Right)* Map of example indoor LOS environment.

and throughput capacity of the central controller along with the frame length. For example, when the central controller becomes the bottleneck, the frame length can be increased to support more antennas, or measurement can be easily parallelized across more servers. Since the central controller is implemented modularly and uses a standard Ethernet interface, it is straightforward to port other Software-Defined Radio (SDR) hardware to this measurement system.

In the results reported in this paper, users were configured to send 802.11 Long Training Symbols (LTS) pilots, which have 52 subcarriers, in a Time-Division Multiple Access (TDMA) fashion at the beginning of each frame, however the system flexibly supports arbitrary pilot formats. This configuration enables wideband pilots for 8 users, along with a noise estimate, to be collected in less than 100 μs.

More information and updates on the measurement tools can be found online [8].

A. Continuous Measurements

In order to achieve the highest possible time-frequency resolution, as limited by the hardware sample rate, we also developed a WARPLab [24] based measurement system that is able to take continuous time measurements for up to 6 s at 20 MHz or 3 s at 40 MHz. However, this measurement system does not implement realtime continuous AGC and can only support one user. In this setup, continuous IQ samples are saved locally to memory at each WARP node, then offloaded to the central controller after the trace is complete. These traces are saved in the same format as the Python framework, enabling the same analysis tools to be applied to both types of traces. This system was developed to provide insight in to how user CSI behaves in real world environments at a sample-level resolution.

IV. Measurement Campaign

We conducted an extensive measurement campaign that includes fully mobile traces across the UHF, 2.4 GHz, and 5 GHz bands in diverse environments. At 2.4 and 5 GHz, we collected traces with up to 104 base-station antennas serving 8 users; at UHF, we collected traces with up to 8 base-station antennas serving 6 users. We collected measurements both indoor and outdoor environments, with varying mobility. These traces typically have frame lengths, i.e., time resolution, varying from 2 ms to 50 ms.

Our primary measurement campaign for this paper consists of 96 base-station antennas serving 8 users at 2.4 and 5 GHz, and thus our analysis will focus on those measurements, particularly since, as shown by Fig. 11, comparing traces of different dimensions is not always equitable. The core of these measurements consists of 3 topologies, indoor NLOS, indoor LOS, and outdoor LOS, at 2.4 GHz and 5 GHz, with five levels of mobility: stationary, environmental mobility, user mobility, linear track mobility, and 'naturalistic' mobility. Stationary traces were taken at night with little or no mobility either at the users or in the environment. Environmental mobility consisted of two people intentionally walking around the mobiles and opening and closing doors. User mobility consisted of two people physically picking up the battery-powered mobiles and walking around with them, while other users were stationary. For controlled motion, we used a CineMoco track [27] to linearly move one user, with other users stationary and limited or no other environmental mobility. To simulate somewhat realistic usage, we also conducted a 'naturalistic' scenario where a person picked up one of the user devices and pretended it was a cell phone, then set it back down, with all other users stationary. Additional traces with controlled rotational mobility, two-antenna users, and other test-specific setups were also taken in some topologies. In most traces, AGC was set then disabled to avoid gain jumps during the trace, however this setting, including gain settings, are recorded in each trace.

Unless otherwise specified, these measurements used omni-directional monopole antennas arranged in an 8 by 12 array spaced 63.5 mm apart, which is half a wavelength at 2.4 GHz, and approximately 1 wavelength at 5 GHz.

For analysis purposes, we found collecting relatively short traces of 20 to 120 s with a 2 ms channel resounding interval worked best for limiting the amount of data and processing time, while enabling us to investigate the behavior of specific scenarios. However, to investigate the longitudinal behavior of users in an office environment, we collected a few traces lasting 20 minutes with reduced time resolution.

Every measurement topology is well documented with experimental descriptions, maps, photos, and, in some cases, even video, which are included with each trace online. Examples of the different environments are shown in Fig. 2. More information on the UHF measurements can be found in [28]. To date, the online repository already contains over 100 traces spanning 20 topologies, providing over one billion channel measurements and 1 TB of data.

V. Channel Analysis Toolbox

To enable rapid analysis of these traces we provide a modular Python-based analysis toolbox, which leverages NumPy and SciPy [29] to accelerate computation. The toolbox converts the raw IQ samples in trace files to CSI, then is able to efficiently compute many useful channel characteristics, including correlation, coherence, Demmel condition number, achievable capacity, impact of delayed CSI, and more. All results presented in this paper were computed using this channel analysis toolbox, which is made available open-source online [8]. In the analysis reported in this paper we are primarily concerned with the following channel characteristics:

Correlation is the normalized product of a user's 1-dimensional CSI, $\mathbf{h}_i$, with the conjugate of a user's CSI, $\mathbf{h}_j$: $\delta(t) = \frac{|\mathbf{h}_j^H(t)\mathbf{h}_i(t_1)|}{\|\mathbf{h}_j(t)\|\|\mathbf{h}_i(t_1)\|}$. Correlating the same user's CSI, $i = j$, at different points in time provides insight in to the channel stability; this auto-correlation corresponds to the square root of the received power if the base station were to serve that user with single-user beamforming. The cross-correlation of one user's CSI with another user's CSI provides insight in to the channel orthogonality and inter-user interference.

Coherence is the statistically expected auto-correlation of a channel with itself given a time delay of Δ: $\rho_i(\Delta) = \mathbb{E}[\frac{|\mathbf{h}_i^H(t-\Delta)\mathbf{h}_i(t)|}{\|\mathbf{h}_i(t-\Delta)\|\|\mathbf{h}_i(t)\|}]$. Channel coherence provides statistical insight in to the expected behavior of channels across time, and is useful for comparing different mobilities, propagation environments, and frequencies. Notably, some previous work, e.g., [11], [15], [16], considers each SISO channel separately, thus discarding the relative phase information between base-station antennas. In this paper we use the expected auto-correlation of the MIMO channel, that is, the auto-correlation of the user's entire CSI vector, which provides much better insight in to the performance of MU-MIMO. *Coherence time* is defined as the delay before the expected correlation falls below a certain threshold, typically 0.95 or 0.90.

Demmel condition number, $d \in [K, +\infty)$ is the ratio of the sum of the eigenvalues to the minimum eigenvalue: $d = \frac{\sum_{k=1}^{n} \lambda_k}{\lambda_n}$, where $\lambda_1 > \lambda_2 > \ldots > \lambda_n$ are the eigenvalues of the matrix $\mathbf{H}\mathbf{H}^H$ [30]. The Demmel condition number is a key indicator of MU-MIMO performance for a given set of users. In multi-user conjugate beamforming systems, it indicates the amount of inter-user interference to be expected; in zeroforcing, it indicates the amount of power that has to be sacrificed to form the nulls to each user.

Capacity, or achievable rate, is computed as the aggregate Shannon capacity of each user's MU-MIMO channel: $C(t) = \sum_{i=1}^{K} log_2(1 + \frac{|\mathbf{w}_i(t-\Delta)\mathbf{h}_i^T(t)|^2 P_i}{\sum_{j \neq i} |\mathbf{w}_j(t-\Delta)\mathbf{h}_i^T(t)|^2 P_j + N})$, where $\mathbf{W}$ are the beamweights for a linear beamformer, and P is the per-user transmission power. For multi-user conjugate $\mathbf{W}$ is simply the complex conjugate of the channel matrix, $c\mathbf{H}^*$, where c is a power scaler. For zeroforcing, $\mathbf{W}$ is the psuedo-inverse of the channel matrix, $c\mathbf{H}^*(\mathbf{H}^T\mathbf{H}^*)^{-1}$. It is important to note that beamforming weights should never be applied to the CSI they were derived from to determine capacity, as not only is this impossible in a real system, but also because it correlates noise power, resulting in inaccurate capacity. To avoid this, our capacity analysis uses the first LTS pilot symbol to derive the beamweights, then applies them to the second LTS pilot symbol; this provides consistency for any Δ, while enabling accurate analysis for $\Delta = 0$. Thus, we can emulate the achievable capacity offline using only recorded CSI.

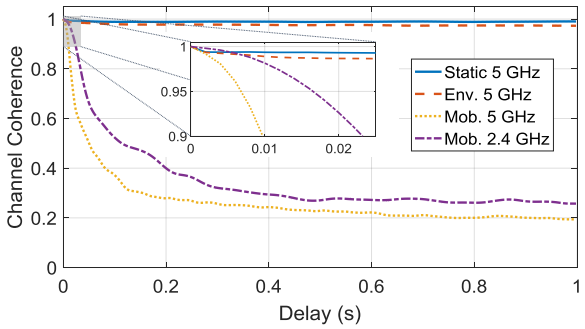

Fig. 3. Channel coherence of 5 GHz for completely stationary topologies (Static), topologies with environmental mobility (Env.), and topologies with the users moving at pedestrian speeds (Mob.) in NLOS environments. Stationary and environmental mobility topologies are long-term stable, whereas mobility drastically reduces channel coherence. To demonstrate the impact of frequency, 2.4 GHz with user mobility is also shown, which has a higher channel coherence than 5 GHz with similar mobility.

Expected capacity is the statistically expected ratio of the achievable rate between when CSI is estimated and when that CSI is used for MU-MIMO given some time delay, Δ: $\gamma(\Delta) = \mathbb{E}[\frac{C(t-\Delta)}{C(t)}]$. Expected capacity can provide insight in to how a MU-MIMO system will perform given a channel resounding interval (Δ) under different mobilities, propagation environments, and frequencies.

VI. Results

We analyze the channel traces from the measurement campaign with regard to the impact of mobility, environment, and base station scale on the channel stability and capacity. Many of these results are intuitive and expected, however our measurements characterize ground-truth points for real-world environments.

A. Impact of Mobility

We investigate six levels of mobility:

1) Stationary Environments and Users: Stationary environments are long-term stable over periods of tens of minutes, with no indication of changing, regardless of carrier frequency. As Figs. 3 and 4 show, stationary channels have a virtually indefinite channel coherence. In many other traces we have observed channels in stationary environments be stable for hours, maintaing a channel correlation above 0.98.

2) Environmental Mobility: As Fig. 3 shows, environmental mobility slightly reduces the coherence of the channel, however it still stays above 0.97 for over a second. Longitudinal studies of the channels in a real-world office environment show similar results; we see in Fig. 4 that channels are seemingly indefinitely stable other than brief periods where the channel is altered, e.g., when a person walks between the user and base station, occluding some of the paths. Typically we observe

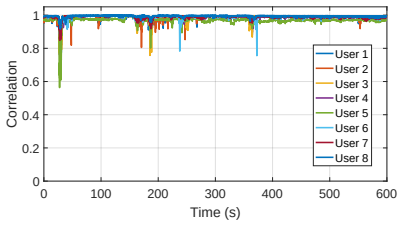

Fig. 4. Each users' auto-correlation with a single channel measurement at 300 s in NLOS at 2.4 GHz. Even with environmental mobility stationary users' channels are remarkably stable: other than brief interruptions, their channel correlation typically stays above 0.95.

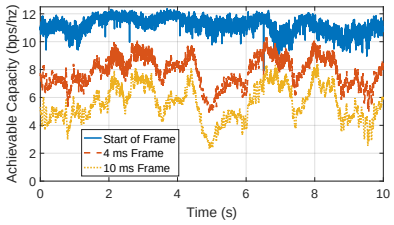

Fig. 5. Achievable capacity at start of a frame, i.e., immediately after CSI collection, as well as at the end of either a 4 ms or 10 ms frame, for a single user with pedestrian mobility in a 96x8 zeroforcing system at 2.4 GHz.

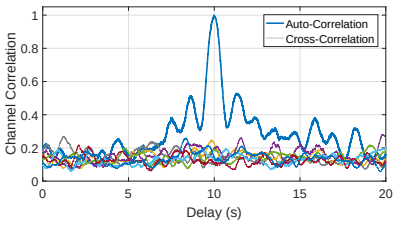

Fig. 6. Correlation with user at 10 s on a linear track moving at 4̃.6 cm/s in LOS at 5 GHz. The top (blue) curve is the moving user's auto-correlation, the lower curves are the cross-correlation with the other 7 users. The 20 s shown corresponds to approximately 92 cm.

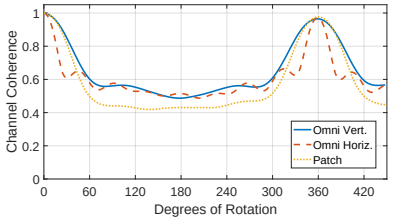

Fig. 7. Channel coherence of a user with controlled rotational mobility at 2.4 GHz in NLOS using an omnidirectional antenna and a patch antenna. Rotational mobility can also significantly impact channel stability; the antenna's non-isotropic radiation pattern creates additional spatial selectivity.

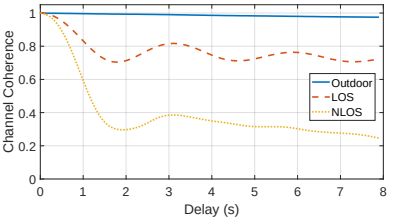

Fig. 8. Channel coherence at 2.4 GHz on a linear track moving at 4.6 cm/s. NLOS environments are less stable than LOS environments with mobility, and outdoor environments with little multipath are very stable. The 8 s shown corresponds to approximately 37 cm of movement. Note that the sinusoidal behavior corresponds with a wavelength, e.g., an interference pattern from reflections.

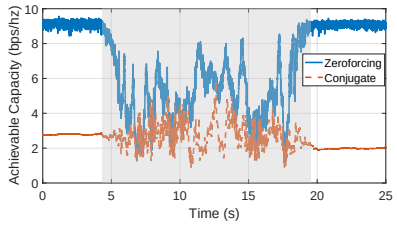

Fig. 9. Capacity plot for a single user with naturalistic mobilitiy, e.g., a cell phone being picked up and used in a 96x8 system with a 10 ms resounding interval. The shaded region indicates when the user was moving. We see the capacity is very bimodal with the channel stable while the user is stationary, and very unstable while the user is moving.

that each user experiences independent fades, however in some cases mobility near the base station can cause fades for multiple users, as is the case at approximately 30 s in Fig. 4. Thus the capacity of many-antenna MU-MIMO systems with environmental mobility is still very stable, but can exhibit brief drops of up to 45%. Notably, the frequency and severity of these fades is quite clearly dependent on the topology and actual mobility.

3) User Mobility: User mobility drastically impacts channel stability, as it affects all paths across the entire frequency band. As shown in Fig. 3, with user mobility 0.90 channel coherence drops to 9.5 ms and 23 ms for 2.4 GHz and 5 GHz, respectively, in a 96x8 MU-MIMO system with NLOS propagation. Similarly, 0.95 channel coherence is 7 ms and 16 ms, respectively. This has a huge impact on achievable capacity, as a system that estimates the user channels every 10 ms will lose an average over 50% of capacity, with dropouts as high as 80%, even with just pedestrian mobility at 5 GHz, as shown in Fig. 5.

4) Track Mobility: By controlling mobility with a linear track we are able to draw interesting insight in to the interference patterns in different environments and at different frequencies. Fig. 6 demonstrates the correlation of the user on the track with itself across time, and thus space, as well as its cross-correlation with the other 7 users, in a 5 GHz NLOS environment. This allows us to visualize how the intended signal strength and inter-user interference varies across space. For example, if the system were to use multi-user conjugate, the top (blue) curve represents the amplitude of the intended signal for the user on the track, whereas the other curves each represent the interference from other users. Similarly, this same method can be used to visualize the signal strength and interference of specific beamforming techniques, such as zeroforcing.

5) Rotational Mobility: We found rotational mobility, where the user's antenna rotates in space, to have a significant impact on channel stability. Leveraging a controlled stepper motor we rotated a user's antenna at a constant rate of 18.75 rpm, then measured the channel at 2.4 GHz NLOS to a 96-antenna base station. Due to the connector and mounting, the rotating antenna was offset 2 cm from the center of rotation. Using this experimental setup we collected traces with the following user antennas: (*i*) a 14 cm omnidirectional 3 dBi monopole oriented vertically; (*ii*) the same antenna oriented horizontally, its tip thus having a radius of 16 cm from the center of rotation; and (*iii*) a 6 dBi patch antenna with an 80 degree beamwidth in both azimuth and elevation. As shown in Fig. 7, within 25 degrees of rotation, all three scenarios fall below below 0.90 channel coherence, with the horizontally mounted antenna reaching 0.90 coherence with just 12.5 degrees of rotation. Unsurprisingly, when the antennas reach their original position, at 360 degrees, we see coherence returns to over 0.97, as should be expected. Rotational movement occurs in the pedestrian and naturalistic traces, e.g., Figs. 3 and 9, and thus likely contributes to the low channel stability. As rotational mobility is known to be

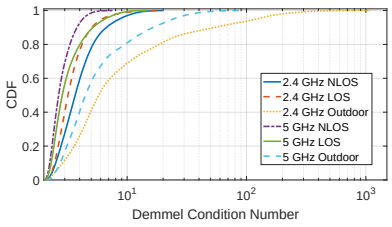

Fig. 10. Cumulative distribution function (CDF) of average Demmel condition number for NLOS, LOS, and outdoor propagation environments at 2.4 GHz and 5 GHz in a 96x2 system. In outdoor environments the lack of multipath can make it much harder for MU-MIMO to separate users.

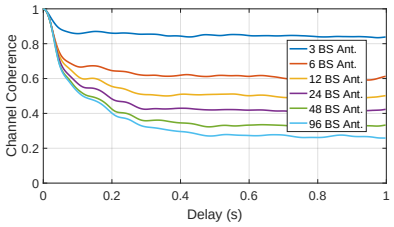

Fig. 11. Channel coherence vs. number of base-station antennas with pedestrian mobility in 2.4 GHz NLOS. Scaling up the number of base-station antennas significantly reduces coherence.

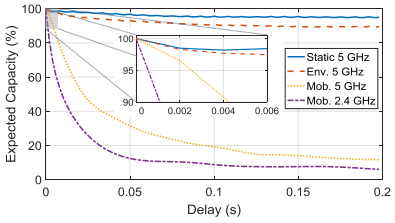

Fig. 12. The expected achievable capacity of a mobile user in a 96x8 MU-MIMO system vs. channel resounding interval in NLOS. The channel resounding interval required to maintain 90% average capacity can be an order magnitude lower than the measured channel coherence.

common in mobile device usage [31], our result suggests that it should be considered in MU-MIMO system design and modeling.

6) Naturalistic Mobility: We find that realistic mobility, e.g., the user equipment is held and moved to mimic cell phone use, results in bimodal channel stability; that is, the user is typically either stationary or moving, which results in very stable and very unstable channels, respectively. As shown in Fig. 9, we see that initially the channel is very stationary, but at about 5 s the user receives a call and begins to move around until about 20 s when the call ends and the user sits down, causing the channel to stabilize. A video of this mobility is included with the trace online, and there are other similar traces in other topologies.

B. Impact of Environment

The environment has a strong impact on user orthogonality and channel stability.

As expected, more multipath increases user orthogonality, and therefore capacity, as MU-MIMO can more easily separate multiple users. As shown in Fig. 10, the lack of multipath in outdoor propagation environments makes the channel poorly conditioned, meaning that it is much harder for MU-MIMO to separate users efficiently. This indicates that MU-MIMO can strongly benefit from multipath, and that user selection in high multipath environments is less critical than environments with low multipath.

However, this better channel condition with multipath comes at a cost: the additional spatial selectivity makes the channels much less stable in the presence of user mobility, as shown by Fig. 8. In outdoor environments the beamwidth created by the base station's antenna aperture dominates channel coherence, whereas in high multipath NLOS environments the channel coherence is dominated by the motion relative to the wavelength. Notably, in our experiments with the track spaced approximately 7 m in front of the array, with a horizontal aperture size of 0.5 m, both 2.4 and 5 GHz have horizontal beamwidths of more than 1.25 m, thus the coherence in the outdoor measurement is expected.

While it is difficult to quantize generally, based on our topologies we also found that NLOS environments are typically not as affected by environmental mobility, since, due to the spatial diversity, it is rare for environmental movement to affect a large portion of paths and frequencies. In contrast, LOS environments show a more bimodal behavior: environmental movement in the direct path causes more extreme changes in the channel, whereas movement not in the direct path rarely has much effect. Thus, with stationary users and environmental mobility, LOS topologies exhibit deep fades more rarely, whereas NLOS topologies exhibit less severe fades more frequently.

C. Impact of Increasing Number of Base-Station Antennas

In Fig. 11 we see that increasing the number of base-station antennas significantly decreases coherence time. More antennas introduce more spatial selectivity, i.e., a narrower beam width, thus movement has a stronger impact on the channel. This effect is mathematically expected and fundamental to [6]; as the number of base-station antennas increases, different spatial locations become more orthogonal. Notably, we found this reduction in channel coherence does not necessarily reduce capacity, as users can actually be moving away from interference from other users, and additional antennas help suppress inter-user interference.

D. System Implications

These channel characteristics have important implications for many-antenna MU-MIMO system design. Due to the instability of capacity under just pedestrian mobility, e.g., losses of up to 50% within 4 ms of channel sounding, shown in Fig. 5, selecting the Modulation and Coding Scheme (MCS) for users will be critical, and likely necessitates fine-grain adaptive or rateless coding schemes, such as [32], [33]. The significant differences we observed in channel coherence and orthogonality across environments and frequencies also must be considered in MU-MIMO systems, particularly for user selection algorithms and CSI collection. Furthermore, the bimodal mobility inherent in naturalistic user movement, shown in Fig. 9, can be leveraged to significantly reduce channel sounding overhead.

We also find that measured channel coherence is *not* an accurate estimate of the channel resounding interval in a MU-MIMO system, which is corroborated in [13], [34]. The 90% expected capacity, shown in Fig. 12, for 2.4 and 5 GHz are 1.1 ms and 4.2 ms, respectively. For 2.4 GHz, this is over 20 times lower than the 0.9 coherence time for the same trace.

We counterintuitively observe that the expected capacity drops more quickly for 2.4 GHz than 5 GHz in this scenario. This is because capacity is dependent on a many factors, including the beamformer, the number of other users, their orthogonality, and the environment. In particular, the initial beamformed channel SINR significantly affects expected capacity, as high SINR channels are inherently more unstable; this is actually the predominant reason for the rapid degradation of 2.4 GHz capacity in Fig. 12. To be clear, channel coherence can be defined, and measured, in many ways, e.g., the theoretical block-fading model assumes the channel does not change during a coherence time interval. Thus using channel coherence interchangeably with the channel resounding interval is not necessarily incorrect, however most *measurements* of channel coherence do not account for all of the factors that affect system capacity, which certainly should be considered in the design of channel sounding protocols for MU-MIMO.

VII. Conclusion

To achieve their full potential, next-generation many-antenna MU-MIMO wireless systems need to carefully consider real-world channel properties, particularly mobility. To help guide their design, we implemented a realtime wideband many-antenna MU-MIMO channel measurement system capable of taking high time-frequency resolution traces in UHF, 2.4 GHz, and 5 GHz bands. Leveraging this platform we took a comprehensive set of reference measurements in diverse environments with varying mobility. Our fundamental analysis of these channels regarding stability and capacity provides critical insight in to the factors that impact MU-MIMO performance, and highlights the challenges mobility presents to many-antenna MU-MIMO system design.

These measurement tools, channel traces, and analysis toolbox are made freely available online [8]. We intend these traces and tools to help guide the design of next-generation MU-MIMO systems.

Acknowledgments

This work was supported in part by NSF grants EARS 1444056, CRI 1405937, CNS 1518916, SBIR 1520496, and SBIR 1632565. We thank Songtao He, Abeer Javed, Narendra Anand, Rahman Doost-Mohammady, Josh Blum, Ashutosh Sabharwal, and Edward Knightly for their input, support, and help. We appreciate the support of the Xilinx University Program. Finally, we thank Christoph Studer for the invitation to present this work at Asilomar.

References

[1] 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*, 2012.

[2] E. Larsson, O. Edfors, F. Tufvesson, and T. Marzetta, "Massive MIMO for next generation wireless systems," *IEEE Communications Magazine*, 2014.

[3] N. Choubey and A. Panah, "Introducing Facebook's new terrestrial connectivity systems - Terragraph and Project ARIES," https://code.facebook.com/posts/1072680049445290, 2016.

[4] C. Shepard, H. Yu, and L. Zhong, "ArgosV2: A flexible many-antenna research platform," in *Extended Demonstration Abstract in Proc. ACM MobiCom*, 2013.

[5] C. Shepard, A. Javed, and L. Zhong, "Control channel design for many-antenna MU-MIMO," in *Proc. ACM MobiCom*, 2015.

[6] T. L. Marzetta, "Noncooperative cellular wireless with unlimited numbers of base station antennas," 2010.

[7] B. Hassibi and B. M. Hochwald, "How much training is needed in multiple-antenna wireless links?" *IEEE Transactions on Information Theory*, vol. 49, no. 4, pp. 951–963, 2003.

[8] The Argos project website, http://projectargos.org.

[9] Samsung and Nokia Networks, "New SID proposal: Study on elevation beamforming/full-dimension (FD) MIMO for LTE," *http://www.3gpp.org/ftp/tsg_ran/tsg_ran/TSGR_65/Docs/RP-141644.zip*, September 2014.

[10] C. Shepard, N. Anand, and L. Zhong, "Practical Performance of MU-MIMO Precoding in Many-antenna Base Stations," in *Proc. ACM CellNet Workshop*, 2013.

[11] K. Nishimori, Y. Takatori, and W. Yamada, "Measured Doppler Frequency in Indoor Office Environment," IEEE, 2009.

[12] G. Breit, "Coherence Time Measurement for TGac Channel Model," IEEE, 2009.

[13] R. Kudo, K. Ishihara, and Y. Takatori, "Measured channel variation and coherence time in NTT lab," *IEEE 802.11-10/0087r0*, 2010.

[14] N. Honma, K. Nishimori, Y. Takatori, and W. Yamada, "Effect of SDMA in 802.11ac ," IEEE, 2009.

[15] J. Wallace, M. A. Jensen, A. L. Swindlehurst, and B. D. Jeffs, "Experimental characterization of the MIMO wireless channel: Data acquisition and analysis," *IEEE Transactions on Wireless Communications*, 2003.

[16] N. Anand, R. E. Guerra, and E. W. Knightly, "The case for UHF-band MU-MIMO," in *Proc. ACM MobiCom*, 2014, pp. 29–40.

[17] S. Yoo, S. Kim, Y. Son, J. Yi, and S. Choi, "Practical antenna selection for wlan ap," in *Proc. IEEE INFOCOM*, 2016.

[18] H. Hofstetter, I. Viering, and W. Utschick, "Evaluation of suburban measurements by eigenvalue statistics," in *Proceedings of the 4th COST*, vol. 273. Citeseer, 2002.

[19] X. Gao, F. Tufvesson, O. Edfors, and F. Rusek, "Measured propagation characteristics for very-large MIMO at 2.6 GHz," in *Proc. IEEE Asilomar Conference on Signals, Systems, and Computers*, 2012.

[20] X. Gao, O. Edfors, F. Rusek, and F. Tufvesson, "Massive MIMO performance evaluation based on measured propagation data," *IEEE Transactions on Wireless Communications*, 2015.

[21] J. Hoydis, C. Hoek, T. Wild, and S. Ten Brink, "Channel measurements for large antenna arrays," in *IEEE Int. Symp. Wireless Communication Systems (ISWCS)*, 2012.

[22] À. O. Martínez, E. De Carvalho, and J. Ø. Nielsen, "Towards very large aperture massive mimo: A measurement based study," in *2014 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 2014, pp. 281–286.

[23] J. Flordelis, X. Gao, G. Dahman, F. Rusek, O. Edfors, and F. Tufvesson, "Spatial separation of closely-spaced users in measured massive multi-user MIMO channels," in *Proc. IEEE Int. Conf. Communications (ICC)*, 2015.

[24] Rice University Wireless Open Access Research Platform, http://warpproject.org.

[25] The HDF Group HDF5, https://www.hdfgroup.org/HDF5/.

[26] H5PY HDF5 for Python, http://www.h5py.org/.

[27] Cinetics CineMoco System, https://cinetics.com/kit/cinemoco_system.

[28] R. E. Guerra, N. Anand, C. Shepard, and E. W. Knightly, "Opportunistic channel estimation for implicit 802.11af MU-MIMO," in *Proc. of First International Conf. in Networking Science & Practice (ITC)*, 2016.

[29] Scientific Computing Tools for Python, http://www.numpy.org/.

[30] C. Zhong, M. R. McKay, T. Ratnarajah, and K.-K. Wong, "Distribution of the Demmel condition number of wishart matrices," *IEEE Transactions on Communications*, 2011.

[31] A. Amiri Sani, L. Zhong, and A. Sabharwal, "Directional antenna diversity for mobile devices: characterizations and solutions," in *Proc. ACM MobiCom*, 2010.

[32] A. Svensson, "An introduction to adaptive qam modulation schemes for known and predicted channels," *Proceedings of the IEEE*, vol. 95, no. 12, pp. 2322–2336, 2007.

[33] A. Shokrollahi, "Raptor codes," *IEEE transactions on information theory*, vol. 52, no. 6, pp. 2551–2567, 2006.

[34] X. Xie, X. Zhang, and K. Sundaresan, "Adaptive feedback compression for MIMO networks," in *Proc. ACM MobiCom*, 2013.