Mobility Resilience and Overhead Constrained Adaptation in Directional 60 GHz WLANs: Protocol Design and System Implementation

Muhammad Kumail Haider and Edward W. Knightly
Rice University, Houston, Texas, USA
{kumail.haider,knightly}@rice.edu

ABSTRACT

High directivity of 60 GHz links introduces new link training and adaptation challenges due to both client and environmental mobility. In this paper, we design, implement and evaluate **MOCA**, a protocol for **M**obility resilience and **O**verhead **C**onstrained **A**daptation for directional 60 GHz links. Since mobility-induced link blockage and misalignment cannot be countered with data rate adaptation alone, we introduce *Beam Sounding* as a mechanism invoked before each data transmission to estimate the link quality for selected beams, and identify and adapt to link impairments. We devise proactive techniques to restore broken directional links with low overhead and design a mechanism to jointly adapt beamwidth and data rate, targeting throughput maximization that incorporates data rate, overhead for beam alignment, and mobility resilience. We implement a programmable node and testbed using software defined radios with commercial 60 GHz transceivers, and conduct an extensive over-the-air measurement study to collect channel traces for various environments. Based on trace based emulations and the IEEE 802.11ad channel model, we evaluate MOCA under a variety of propagation environments and mobility scenarios. Our experiments show that MOCA achieves up to $2\times$ throughput gains compared to a baseline WLAN scheme in a diverse set of operational conditions.

Categories and Subject Descriptors

C.2.1 [**COMPUTER-COMMUNICATION NETWORKS**]: Network Architecture and Design—*Wireless communication*

Keywords

60 GHz, IEEE 802.11ad, Beamforming, Protocol Design, Mobility Resilience, Adaptation, Measurement, software-radio

1. INTRODUCTION

The IEEE 802.11ad standard for 60 GHz WLANs defines specifications to achieve up to 6.7 Gbps rates by using high gain phased array antennas [4, 9]. By adjusting weights of the antenna elements via a predefined beamforming codebook, beam patterns (or virtual "sectors") of specific beamwidth and direction can be achieved. Since both the transmitter (Tx) and the receiver (Rx) use directional beams, the standard also defines a Beamforming Training (BFT) procedure that uses exhaustive search to discover the Tx-Rx pair of sectors with the highest signal strength. Unfortunately, such directional links can be impaired by *(i) blockage* due to mobile obstructions (e.g., 15 to 20 dB loss for the human body [16]); and *(ii) misalignment* of the Tx and Rx sectors due to rotational or translational mobility of nodes. Without adaptation, both scenarios would lead to degradation of the link budget associated with the selected sectors, and the signal strength may drop below the minimum sensitivity threshold, thus rendering the selected sector pair unusable, i.e., breaking the link. In 802.11ad, such mobility can yield repeatedly broken links, thereby requiring BFT and incurring the overhead associated with exhaustive training each time.

In this paper, we design, implement, and evaluate Mobility resilience and Overhead Constrained Adaptation (MOCA), a protocol for highly directional 60 GHz WLANs. Our objective is to maximize throughput incorporating antenna sector selection and alignment, PHY rate adaptation, beamwidth selection, and training overhead. In particular, we make the following contributions.

First, we devise *Beam Sounding*, a short control-frame exchange prior to data transmissions with two objectives: *(i)* We use channel estimation and training fields to assess SNR for the current sector pairs, perform beam refinement to fine tune antenna weights for improved directivity gain, and select the highest supportable PHY rate. *(ii)* In case the exchange fails, the transmitter infers blockage or misalignment without data loss thereby enabling rapid identification of link breakage resulting from mobility induced blockage or sector misalignment. In order to avoid 802.11ad's exhaustive-search overhead and any interruption to the established connection in this case, we design *Preemptive Fast Recovery* to proactively search for alternate fail-over sectors in advance, which can potentially be used upon link breakage.

Second, we develop an algorithm to dynamically select and adapt beamwidth in response to nodal and environmental mobility, with an objective to maximize average link throughput. To achieve this, we employ a multi-level codebook design [7, 22], in which each level corresponds to a certain beamwidth, whereas codes at a level constitute antenna element weights to define beams in different directions, spanning the entire 360° at each level. Our key technique is to incorporate both directivity gain and training overhead: For instantaneous data rate maximization, the narrowest beamwidth sectors should be selected since they offer maximum directivity gain. However, in the presence of mobility, the repeated training overhead for the narrowest beams may overwhelm the advantage of higher directivity gain. This overhead results from higher susceptibility of narrower sectors to misalignment, requiring BFT more fre-

MobiHoc'16, July 04-08, 2016, Paderborn, Germany

DOI: http://dx.doi.org/10.1145/2942358.2942380

quently, and the fact that narrow beamwidth sectors require a longer time for BFT since a greater number of sectors need to be searched exhaustively at both ends. Consequently, wider sectors with non-maximum signal strength may result in higher net throughput due to greater resilience and lower overhead. We devise a throughput estimate to capture this rate vs. mobility resilience tradeoff incorporating the maximum rate currently supported on the link, training overhead as a function of beamwidth, and the extent of mobility inferred via frame and Beam Sounding analysis. We predict throughput prior to each channel access for the entire set of available beamwidths, and select the beamwidth and corresponding rate that maximize the predicted average link throughput.

Finally, we develop a 60 GHz programmable testbed using VubIQ 60 GHz transceivers (802.11ad compliant radios) with 1.8 GHz bandwidth and WARP baseband. We conduct an extensive measurement study to collect over-the-air channel traces for a lab, office and classroom. For performance evaluation under various environments and mobility scenarios, we design a custom simulator that encompasses the trace based channel measurements collected in the hardware study, as well as the channel model presented in the 802.11ad evaluation methodology [12]. The latter is an analytical model and includes electronically steerable beam arrays with codebook based beamforming and flexible beamwidths. Our experiments show that compared to exhaustive training, MOCA's Beam Sounding and Preemptive Fast Recovery mechanisms introduce minimal overhead, which is far outweighed by the throughput gains from quick identification of breakage and restoration of links without losing data transmission opportunities. We compare MOCA with a baseline WLAN scheme and find that MOCA achieves up to $2\times$ throughput gains in most blockage and misalignment scenarios.

The remainder of this paper is organized as follows. A brief background on 802.11ad is presented in Sec. 2. Sec. 3 describes the design of various components of MOCA. We present our testbed, measurements and experiments in Section 4. Sec. 5 discusses the related work. We summarize our findings in Sec. 6.

2. BACKGROUND: IEEE 802.11AD

In this section, we present a brief overview of the PHY and MAC layer specifications of IEEE 802.11ad [4, 9].

2.1 Beamforming Training (BFT)

In contrast to IEEE 802.11n/ac techniques to support multiple spatial streams with multiple RF chains [2], 802.11ad requires only a single RF chain for the antenna array. In particular, 802.11ad realizes Tx and Rx beamforming via analog phase shifters that are configured according to a predefined codebook of beamforming coefficients, with each codebook entry referred to as a sector. The standard's BFT process discovers the maximum signal strength Tx and Rx sectors between a pair of nodes (typically the AP and a STA). As illustrated in Fig. 1, BFT comprises of two phases: a mandatory Sector Level Sweep phase, and an optional Beam Refinement Phase. During the sector sweep, the initiator exhaustively switches across all available sectors, transmitting training frames marked with respective sector identifiers. The responder receives in quasi-omni pattern and identifies the highest strength Tx sector. This process, when repeated at both Tx and Rx, yields the Tx/Rx sector pair having the highest signal quality. Because this process is coarse grained, the Beam Refinement Phase can be used to further align the selected sectors to improve directivity gain.

2.2 MAC Layer Specifications

802.11ad divides time into Beacon Intervals, with a structure depicted in Figure 2. This structure specifies periods for beacon transmission, beamforming, management frame exchanges and data transmission slots. The network can be configured for both scheduled access, in which data slots are pre-assigned to stations, and random access, in which nodes can compete for contention based periods using modified 802.11 random access procedures. Moreover, if a random access period is announced in the Data Transmission Interval, nodes use directional transmission and reception to minimize interference and collisions. Therefore, in an idle state, a node senses the channel in pseudo-omni mode until it receives an RTS or data packet preamble. It can then switch to directional reception, using the sectors selected in BFT, to improve link budget.

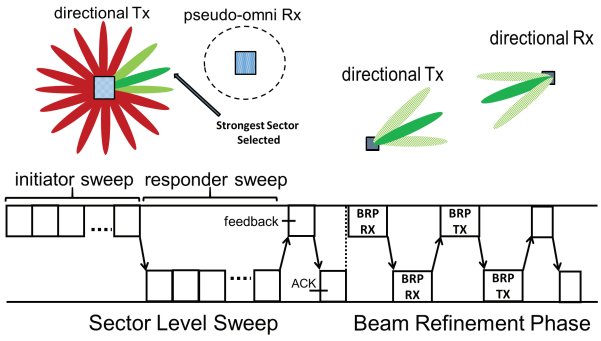

Figure 1: Two phases of Beamforming Training (BFT) 1) Sector Level Sweep in which one node sweeps across all its Tx sectors, while the other receives in pseudo-omni mode and identifies the best Tx sector. 2) Beam Refinement Phase for fine grained calibration of Tx and Rx sectors.

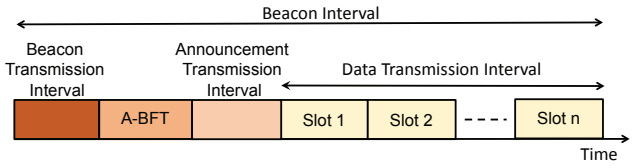

Figure 2: Phases of channel access in a Beacon Interval, starting with a sector sweep of beacon frames by AP. An optional A-BFT slot for beamforming during station association follows, succeeded by a slot for management frame exchange. This is followed by a Data Transmission Interval, which is further divided into (possibly multiple) slots of scheduled and random access.

3. MOCA PROTOCOL DESIGN

In this section, we describe the key components of MOCA to realize mobility resilience and overhead constrained adaptation in 60 GHz WLANs.

3.1 Design Overview

In 802.11ad, Beamforming Training identifies the Tx/Rx sectors with maximum signal strength. With subsequent environmental or nodal mobility, the channel quality of the link can change, which is typically countered by adapting the Modulation and Coding Scheme (MCS) via techniques such as [3, 5, 13]. If even base rate is not feasible, i.e., the current sector pair does not permit communication and the link is broken, BFT is required again to restore the link with a new sector pair. Unfortunately, we will show that this approach of MCS adaptation and BFT yields poor MAC-layer throughput for mobile clients due to the air time required by BFT's exhaustive search. In contrast, MOCA targets to avoid link breakage (and hence BFT) with a preemptive strategy while also replacing the "highest gain sector pair" strategy of 802.11ad with a strategy that jointly incorporates link resilience and search overhead

with directivity gain.

Consequently, we first propose Beam Sounding for identifying misalignment and blockage with low overhead, and avoiding packet losses due to broken links. Our key technique is to introduce a control-frame exchange to probe the selected sectors, before initiating any data transmissions. This probing mechanism assesses the existence and strength of selected sectors and enables inference of link breakage without sending data frames. Moreover, if the link is not broken, we use this feedback to further align the selected sectors and select the highest possible MCS. Moreover, we proactively search for alternate fail-over sectors in advance, which can be used upon link breakage identification to restore the link. This mechanism, if successful, recovers from mobility induced breakage without incurring the exhaustive BFT overhead and without loss of data transmission opportunity.

Second, by searching for the sector-pair with the highest signal strength, 802.11ad BFT implicitly targets a narrow beam with high directivity. In contrast, we dynamically adapt beamwidth in response to the frequency of mobility induced link breakage and the resulting training overhead, with an objective to maximize the average link throughput instead of instantaneous data rate. For this, we employ a multi-level codebook with *K* levels, where each level corresponds to beams of a specific beamwidth, such that the beamwidth decreases with increasing level. For an antenna array of *P* elements, a codebook is defined as a $P \times Q$ matrix, where Q is the number of weight vectors generating beams patterns at a certain level. Thus if $\theta_k (k \in \{1, ..., K\})$ is the beamwidth in radians at the *k*th level, $Q = \frac{2\pi}{\theta_k}$. A DFT-based codebook design [1] can be used to achieve uniform gain beams in all directions, with weights defined as:

$$w(p, q) = [\frac{1}{\sqrt{P}} e^{-j2\pi(p-1)(q-1)/Q}] \quad (1)$$

where $p \in 1, ..., P$ and $q \in 1, ..., Q$.

For a multi-level codebook, existing algorithms (e.g., as proposed in [22]) can be used to generate codebooks with different beamwidths at each level. Consequently, we adapt codebook level, and hence beamwidth, to control link budget and robustness. For example, narrow beamwidths for Tx and Rx sectors can be achieved by traversing deeper levels of beamforming codebook, resulting in higher directivity gain and hence higher link budget. Therefore, for instantaneous data rate maximization, narrowest sectors should be selected since they offer maximum link budget and highest rates. However, selecting a narrow beamwidth also increases the search space for the BFT process as more sectors need to be tested, thereby increasing the training overhead. Likewise, narrow beams are less resilient to client mobility further requiring additional instantiations of BFT.

Using 802.11ad timing values for training frames, if BFT is performed at levels l and m ($l, m \in \{1, ..., K\}$) at the Tx and Rx respectively, then the time required for training (without beam refinement) is:

$$\tau(l, m) = [\frac{a}{\theta_l} + \frac{a}{\theta_m} + c]\ \mu s \quad (2)$$

where a and c are protocol constants, with values 116.24 and 71 μs in our implementation. Notice that the time required for beam training varies inversely to the beamwidth, and for deeper codebook levels, it may even become greater than the data transmission slots. For example, for 3° beamwidth (minimum allowed in 802.11ad), BFT requires ~ 5 *ms* while the maximum transmit slot size is 2 *ms*. Moreover, our channel measurement data in Section 4 shows that narrow beamwidth also makes the selected sectors more susceptible to misalignment, as a slight rotation or sideways motion of the node can result in high degradation in directivity gain. More frequent misalignment instances lead to a higher frequency of performing BFT and hence a higher overhead. Wider beamwidth selection, on the other hand, results in sector pairs with more resilience towards mobility, at the cost of lower data rates. Therefore, in MOCA, we devise an algorithm to dynamically adapt beamwidth in response to mobility such that the average link throughput is maximized.

3.2 Beam Sounding

Nodal or environmental mobility can necessitate adjustment of MCS and beamforming weights in order to avoid packet loss. MOCA comprises several mechanisms that enable link adaptation and optimization without requiring packet loss to infer that adaptation is needed, and without requiring overhead-intensive BFT. We introduce Beam Sounding as a short control-frame exchange between the transmitter and the receiver prior to data transmissions. The key technique is to first test that the current selection of sector pairs yields a link which can support at least the base rate MCS. This ensures that the link exists and that data transmission is feasible without transmitting data itself. We design these frames for a request and response feedback exchange, using the beamforming sectors selected by the pair of nodes in the most recent BFT instance. We use a short structure for Beam Sounding frames, including Short Training Field and Channel Estimation Field and transmit the sounding frames directionally, using the sector which needs to be probed, while the receiver is in pseudo-omni reception. Since Beam Sounding is a request and response procedure, both nodes not only test the existence of the link for a certain selection of sectors, but they can also perform beam alignment and channel estimation.

Thus, Beam Sounding frames serve the following purposes: *(i)* to confirm the usability of the previously selected sector pair. If the Beam Sounding frame exchange fails, MOCA initiates Pre-emptive Fast Recovery (described below) to restore the link. *(ii)* to provide inference about failure of selected link due to breakage (misalignment or blockage), which we use to estimate the extent of mobility in the network (also described below); *(iii)* to perform beam alignment using the Short Training Field prior to packet transmissions to fine tune the Tx and Rx sectors, thereby improving directivity gain; *(iv)* to estimate SNR for MCS selection using the Channel Estimation Field. The required overhead for Beam Sounding frames is approximately 10 μs, orders of magnitude lower than typical BFT epochs and orders of magnitude lower than the time required to transmit and recover from a failed data frame. Moreover, for systems with virtual carrier sense on (although it incurs higher overhead due to sector sweeps for RTS/CTS frames), our Beam Sounding mechanism can be implemented by piggy-backing estimation feedback on RTS and CTS control messages with the aforementioned two fields.

Despite using Beam Sounding to assess link existence and strength, packet loss is possible in MOCA due to channel variations, collisions or MCS over-selection. In this case, the receiver will fail to decode the data frame. However, since the packet preamble is always encoded at MCS-0 with the lowest receive threshold, the receiver may detect the packet transmission if the base rate is supported by the channel. In MOCA, we enable receiver feedback in the form of a negative-acknowledgement (NACK), which helps in packet retransmissions as discussed below.

3.3 Preemptive Fast Recovery

In 802.11ad, a failed transmission triggers a combination of BFT, MCS adaptation, and retransmission. In contrast, with MOCA, we

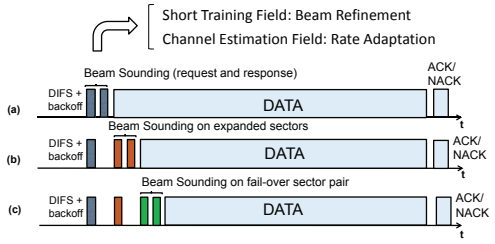

Figure 3: Beam Sounding possibilities in a TXOP in MOCA: (a) Sounding on the primary sector pair is successful, (b) sounding on the primary sector fails, but with expanded sectors, the link is restored; (c) both primary and expanded sectors fail, but a successful exchange of sounding frames occurs for the fail-over sector.

exploit that a failed Beam Sounding exchange (vs. a failed data-ACK exchange) indicates that the current sector pair is not usable. With this early identification that the link is broken, instead of initiating BFT, we introduce a proactive strategy for restoration of the broken link, potentially avoiding the need for exhaustive training. The key idea is to use the short Beam Sounding frames to probe alternate fail-over sector pairs before falling back to BFT. In particular, MOCA retransmits the Beam Sounding frames first by traversing one level higher in the beamforming codebook, yielding a widened sector. Subsequently, if transmission at the new codebook level also fails, MOCA further probes an alternate fail-over sector pair which we discover in advance in the beam training process. Below we first describe the modified training process, and then discuss the complete sequence of frame exchanges.

Resilience Training: The purpose of resilience training is to incorporate a proactive search for alternate fail-over sectors during beam training in advance, which can be used in case of subsequent blockage. If blockage results from environmental mobility, the alternate signal propagation paths, resulting from reflections off of various surfaces, may still exist and can be used opportunistically. We search for any alternate links by modifying the beamforming training procedure such that beam refinement is performed for multiple sector pairs (if they exist), and not just the strongest sector pair. The strongest sector pair from training is selected as the primary link for data transmission, whereas a backup sector pair is stored as an alternate link to be used upon failure of the primary link. It is possible that this recovery will not succeed if the alternate link is also blocked by an obstruction blocking the primary link, or if the alternate link is lost due to environmental or nodal mobility. Nonetheless, since more than 90% of Beamforming Training time comprises of sector sweeps [4], the overhead of MOCA's proactive alternate link search is relatively low.

The sequence of Beam Sounding frame exchanges in MOCA is illustrated in Fig. 3 for different scenarios. As illustrated, MOCA initiates a transmit opportunity (TXOP) by first transmitting sounding frame on the primary sector pair (selected during the latest BFT). If feedback is received, the data packet is transmitted after selecting the maximum MCS, and may result in an ACK or a NACK. In case of a NACK, MOCA implements joint adaptation based on the feedback (described below) and initiates Beam Sounding again in the next TXOP (when channel access is granted).

If the aforementioned sounding fails on the primary sector pair, MOCA infers link breakage due to blockage or sector misalignment and invokes two recovery mechanisms. First, Beam Sounding is performed under new antenna weights obtained by traversing the codebook tree up one level, thereby widening the sector potentially overcoming slight misalignment by using this wider sector. The transmitter also sets a flag in the sounding frame instructing the receiver to expand its sectors as well. If a response is received, the nodes switch to this expanded sector pair as their new primary link, select an MCS based on the feedback, and transmit data.

Second, if this sector expansion fails to restore the link, MOCA uses a second retransmission of sounding frames to probe the alternate fail-over sector pair (at the same codebook level as the original primary sectors) discovered in the resilience training to recover from blockage. If the alternate link is available, the receiver can identify that the probe request corresponds to the fail-over sector based on the mutually agreed sector ID in resilience training. The receiver sends a response on its fail-over sector as well. If this response is received successfully, the nodes switch to this alternate sector pair and proceed to data transmission. This scenario indicates that link breakage was likely due to blockage. In case both proactive strategies fail, MOCA requests the MAC to schedule BFT.

Lastly, we use binary exponential backoff to spread successive probe request attempts to reduce interference and possibility of collisions. However, this backoff counter is separate from the primary backoff counter for channel access and does not impact fairness due to MOCA procedures.

3.4 Joint Adaptation Algorithm

Here, we present the design of MOCA's algorithm to jointly incorporate the beamwidth tradeoff between directivity gain, training overhead, and mobility resilience. We adapt beamwidth to reflect mobility, with wider beams selected to increase resilience with higher mobility, whereas lower mobility favors narrower higher gain beams. Our key technique is to devise an estimate for average link throughput, which captures the rate vs. resilience tradeoff and to drive adaptation from this estimate. We estimate throughput η for beamwidths θ_i corresponding to each level i of the beamforming codebook, and select the beamwidth θ^* that maximizes this estimate.

To derive this estimate, we consider a simplified channel access scenario and only consider a particular sender-receiver pair. We assume that the maximum achievable PHY rate $r(\theta_i)$, for the current beamwidth θ_i is known at the transmitter via BFT and Beam Sounding. If the underlying mobility process is stationary, then the system can be approximated as a renewal process with throughput estimated as

$$\eta(\theta_i) = \frac{(1-\beta(\theta_i))r(\theta_i)t_{slot}}{(1-\beta(\theta_i))(t_{slot}) + \beta(\theta_i)t_{BFT}(\theta_i)} \tag{3}$$

in which t_{slot} is the length of data transmission slot. Moreover, $t_{BFT}(\theta)$ is the time required for BFT for beamwidth θ, capturing the fact that it takes more time to train for narrower beams, even though narrower beams can yield higher $r(\theta)$. We define $\beta(\theta)$ as the probability of a breakage event for beamwidth θ, i.e., the probability of a data transmission being unsuccessful due to link breakage. A successful transmission results in a bit rate of $r(\theta)$ for time t_{slot}, whereas breakage leads to training penalty t_{BFT}. $\beta(\theta)$ has an important role in characterizing the tradeoff between rate and overhead for a particular selection of beamwidth. Under lower mobility conditions, β will also be low, giving more weight to the rate term in the estimate, thus favoring narrower beamwidth. Under higher mobility, the overhead term in the estimate becomes the determining factor favoring wider beamwidths.

A flow diagram for MOCA's decision process is illustrated in Fig.4. At the end of each TXOP, the throughput estimate $\eta(\theta)$ is computed for all available beamwidths based on current estimates of rate $\hat{r}$ and breakage probability $\hat{\beta}(\theta)$. MOCA then selects the beamwidth θ^* that maximizes this estimate. If θ^* is different from

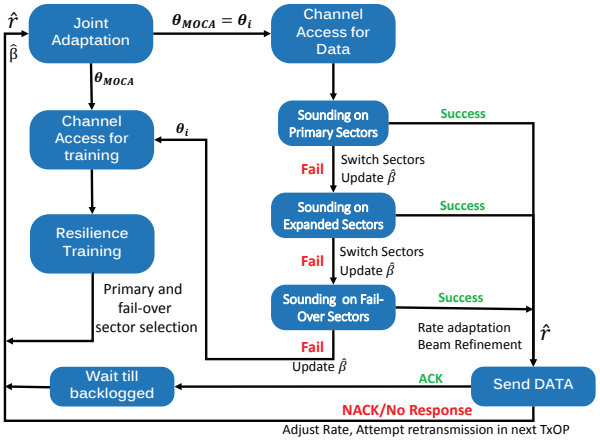

Figure 4: Flow diagram for MOCA

the current beamwidth θ_i, MOCA requests resilience training to select and train sector pairs with beamwidth θ^*. Otherwise, it uses the next TXOP to transmit data by first sending Beam Sounding frames on the primary sector pair since MOCA predicted the current beamwidth to maximize average link throughput. If the sounding frame exchange is successful, MOCA proceeds with packet transmission encoded at rate $\hat{r}$ estimated during Beam Sounding. Otherwise, Beam Sounding frames are used to probe the expanded sectors and the fail-over sectors, as described in Fig. 3. Moreover, in case of a NACK or no response, θ^* is computed again after updating $\hat{\beta}$ and/or $\hat{r}$, and before starting a retransmission.

In MOCA's beamwidth adaptation algorithm, we need an estimate of data rate ($\hat{r}$) and breakage probability ($\hat{\beta}$) for all available beamwidths. The data rate for the current beamwidth $r(\theta_i)$ is already estimated as a result of feedback during BFT and is also refreshed during Beam Sounding frame exchanges. For a different beamwidth θ_j, we estimate $\hat{r}(\theta_j)$ based on the link budget (LB) for θ_i. If the beamwidth used for the Tx and Rx sectors is θ_{Tx} and θ_{Rx} respectively, the LB for a 60 GHz link is approximately:

$$LB(\theta, \phi, d) = G(\theta_{Tx}, \phi) + G(\theta_{Rx}, \phi) + P_{Tx} - PL(d) - \sum_{i=1}^{K} R_i \quad (4)$$

where ϕ_{vert} is the vertical beamwidth and is fixed. G is the antenna gain, P_{Tx} is the transmit power, PL(d) is the path loss at distance d and R_i is the penetration or reflection loss for the i^{th} obstacle. Given that LB(θ_i) is known from the estimation of $\hat{r}(\theta_i)$, we can estimate the change in directivity gain when changing from θ_i to θ_j (since beam refinement is done in Beam Sounding and BFT, we consider directivity gain difference for perfectly aligned beams in this estimate), and account for LB(θ_j), which can then be used to estimate the data rate ($\hat{r}(\theta_j)$) for beamwidth θ_j.

Lastly, the breakage probability β is difficult to directly estimate since it depends on mobility. Consequently, we characterize the extent of mobility via hints from the Beam Sounding mechanism, and estimate β using the recent data and Beam Sounding history, including inferred breakages, in a time period T prior to θ^* computation. We define $\hat{\beta}$ as the fraction of Beam Sounding frames in the time period T, which were lost due to breakage.

In general, β is non-increasing with respect to beamwidth since there is an inverse relation between the width of beamforming sectors and probability of misalignment. For example, doubling the beamwidth will reduce the frequency of misalignment events roughly by half, given the same distance and rotational mobility

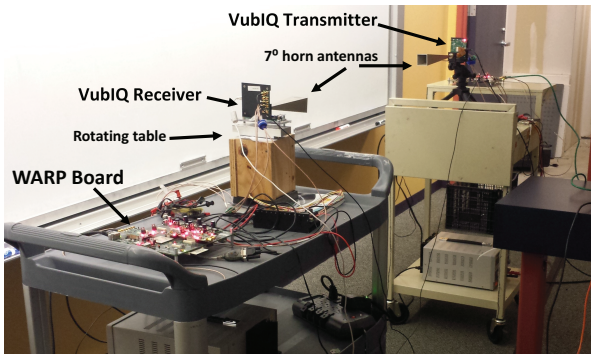

Figure 5: 60 GHz channel measurement experimental setup, with two VubIQ transceivers connected to WARP platforms.

conditions for the node. However, it also depends on the initial alignment of the beams. Since we perform beam alignment in a successful probe exchange, we assume perfect alignment of sectors for θ_i. To capture the impact of this change in beamwidth on β, we use the following approximation, under the perfect alignment assumption:

$$\hat{\beta}(\theta_j) = (\frac{\theta_i}{\theta_j}) * \hat{\beta}(\theta_i) \quad (5)$$

Thus, given the values of $\hat{\beta}(\theta_i)$ and $\hat{r}(\theta_i)$ for the current beamwidth θ_i, MOCA can estimate $\eta(\theta)$ for all $\theta_j \in [\theta_{min}, \theta_{max}]$ and switch beamwidth from θ_i to θ_j if $\eta(\theta_j) > \eta(\theta_i)$. In case a different beamwidth is selected, MOCA initiates channel access for resilience training instead of data transmission.

The choice of interval T for packet history (with n transmission attempts) is a protocol parameter and its choice may make the protocol aggressive or conservative for smaller or larger values respectively. In the model, we assume that the underlying mobility process remains stationary such that the Beam Sounding frames within T yields an estimate of the link breakage probability. This assumption may not hold true if T is large, leading to errors in the estimate. However, as discussed in [18], human blockage and nodal movement events last on the order of seconds, whereas we select T as one beacon interval (100 ms) and assume the underlying mobility process remains stationary as an approximation.

4. IMPLEMENTATION AND EVALUATION

In this section, we first describe our evaluation methodology and the MOCA 60 GHz hardware platform. We then discuss a subset of observations from our channel measurement study for indoor channels using our 60 GHz testbed. Finally, we investigate the performance of MOCA in a variety of environments and mobility scenarios via trace-based experiments using the over-the-air measurements as well as the 802.11ad channel model [12].

4.1 MOCA Implementation

Hardware and Software Platform: We implement all algorithmic aspects of MOCA in software and couple it with the FPGA software-defined-radio WARP combined with a 60 GHz frontend. In particular, we implement a 60 GHz programmable node and testbed using VubIQ [19] transmitter and receiver waveguide system, operating in 57-64 GHz unlicensed frequency band with 1.8 GHz bandwidth (compliant with 802.11ad). We generate I/Q baseband signal at different modulations and rates using WARP baseband [6], and use WARPLab (a framework for rapid physical layer prototyping) to generate BPSK and QPSK baseband signal with 20

MHz bandwidth. The differential I/Q input to the VubIQ transmitter is achieved by feeding WARP's I/Q baseband signal to an evaluation board using a 6 GHz Ultra Dynamic Range Differential Amplifer (ADL5565). To achieve different sector widths, we use horn antennas with 7°, 20°, and 80° beamwidths, and an omni-directional antenna to achieve pseudo-omni reception. Moreover, we achieve mechanical steering by mounting the transceivers on a programmable rotating table which can emulate fixed sectors.

The testbed is shown in Fig. 5. Using this testbed, we conduct over-the-air experiments to collect traces of channel strength and variations for a lab, office and a classroom scenario.

Trace- and Model-Driven Simulation: We developed a custom MATLAB simulator which can accommodate different channel models and mobility scenarios. We use two channel sources to drive the simulator. The first employs over-the-air measurements via the MOCA 60 GHz hardware testbed. The measurements incorporate the effect of path-loss, misalignment of antenna beams, and reflections on signal strength on a 60 GHz channel. To explore a broader set of operational conditions beyond the capabilities of the hardware platform including electronic beam-steering, a larger set of beamwidths, and multiple environmental scenarios, we also implement the channel model used in the 802.11ad channel evaluation methodology [12]. This model covers path loss, channel variations, antenna gains, reflection and penetration losses in a living-room and conference-room environment.

The simulator uses ray tracing and link budget analysis to model signal propagation and detection. We model both environmental and nodal mobility, with nodes capable of translation as well as rotation. For the 802.11ad channel model, the simulator emulates 60 GHz nodes with electronically steerable arrays with flexible beamwidth sectors, based on multi-level beamforming codebook design. The path loss, directivity gain and reflection and penetration losses are dictated by the analytical model. Whereas for the trace-based simulations, we use the same link budget and reflection loss values as measured in the hardware experiments. Therefore, the link budget depends on the distance between the transmitter and the receiver, their relative angles, angle of departure and the angle of arrival (specific values for all these parameters map to a single reading in the measurement data set). Since the measurement data points are relatively coarse-grained (compared to continuous time mobility and rotation in the simulator), we use weighted average to calculate signal strength values at these intermediate points, when using channel traces.

4.2 60 GHz Channel Measurements and System Evaluation

4.2.1 Over-the-Air Measurement Study

Here, we describe a subset of testbed experiments to study propagation at 60 GHz and the impact of mobility and reflected paths.

Line of Sight (LOS) path and angle of arrival: The first over-the-air experiment explores the impact of path loss and angle-of-arrival on signal strength. We measure the Received Signal Strength Indicator (RSSI) and channel variations for a link with an unobstructed LOS path in an 8m x 4m lab. The transmitter is fixed in a corner at 1.5 meters height, whereas the receiver is placed at 1 meter distance from the receiver, at the same height, on a programmable rotating table. We take measurements while the distance between the two nodes is increased to 5 meters in intervals of 1 meter, while maintaining LOS. We also rotate the receiver along the azimuth, from -170° to 170° (with 0° pointing in transmitter direction), and take measurements at intervals of 10 degrees. This rotation results in different angle-of-arrivals at the receiver, depending on receiver orientation. Each measurement consists of 100 over-the-air packet transmissions.

Fig. 6a depicts the variation of normalized RSSI (with respect to maximum value) as the inter-node distance increases from 1-5 meters, for Tx beamwidths 7°, 20°, and 80° (Rx beamwidth is fixed at 20°). 7° beamwidth achieves the maximum signal strength, due to its higher directivity gain, across all distances. At 1 meter, the difference in signal strength between 7° and 20° is 12.2 dB, whereas 20° beamwidth has 11.7 dB higher signal strength than 80°. This difference in directivity gains is within 3 dB of the theoretical gains given by Equation 4. Fig. 6b depicts the received power versus receiver antenna angle for for all three transmit beamwidths. Here the RSSI values are shown for an angular spread of -50° to 50° since there is no signal reception beyond this range. These results illustrate the dominance of the LOS path and presence of weaker reflected components, especially for 7° beamwidth. We also observe that the angular spread of 7° is the largest compared to wider beamwidth sectors. The reason is that the receiver beamwidth is fixed and the transmitter's orientation remains the same for all different transmit beamwidths. This makes the extent of misalignment the same at all receiver angles, and 7° beamwidth shows highest spread due to its much higher signal strength across all angles.

Radial mobility and misalignment: In the next experiment, we investigate the impact of misalignment on link strength. We start with the same set-up as the last experiment, with the Tx and Rx nodes perfectly aligned, and 1m apart. Instead of rotating the receiver, we move it along the circumference of a quad-circle of radius 1m, such that the transmitter is at the center. This results in an offset between beam centers of Tx and Rx antennas (i.e., they are misaligned). The receiver beamwidth in this experiment is 7°. In this experiment, the misalignment results from radial mobility of the receiver around the transmitter. The results in Fig. 6c indicate that the RSSI for 7° beamwidth, though maximum at perfect alignment, decreases sharply as the receiver is moved along the circumference. At relative angles greater than 8° (by extrapolation), signal strength for 20° becomes higher than that for 7°. For 80° transmitter beamwidth, the signal has maximum spread and its strength is higher than lower beamwidths at relative angles greater than 19°.

This experiment also illustrates the impact of beamwidth on link resilience. Although narrower beamwidth achieves higher link strength, it is less resilient to rotation such that a slight misalignment can result in significant degradation of the link strength or may completely break the link. In contrast, wider beams offer lower link budget but provide greater resilience and sufficient signal strength across a larger spread of Tx-Rx relative angles. Moreover, for 80° beamwidth, RSSI starts to increase again after dropping to nearly zero at 63°. This indicates the presence of a reflected path in addition to the LOS path, making the wider beamwidth much more resilient to blockage or misalignment. In this case, the signal spread of narrower beamwidths is not sufficient to exploit this additional path. Therefore, wider beamwidth can also exploit multiple paths in addition to a much wider reception signal along the LOS path, improving resilience.

Reflection experiments: Next, we conduct experiments for the measurement of reflected paths resulting from reflection off of various objects. Here we present an experiment with reflections from a white board in a class-room set-up shown in Fig. 7a. We collect RSSI values for receiver rotation from 0° to 135° in 5° steps, such that at 0° the receiver is pointing towards the transmitter and at 90°, towards the board. The receiver beamwidth is fixed at 7°.

Fig. 7b shows the variation in RSSI vs. receiver angles between 0° and 135° for transmitter beamwidth of 7°, at different distances.

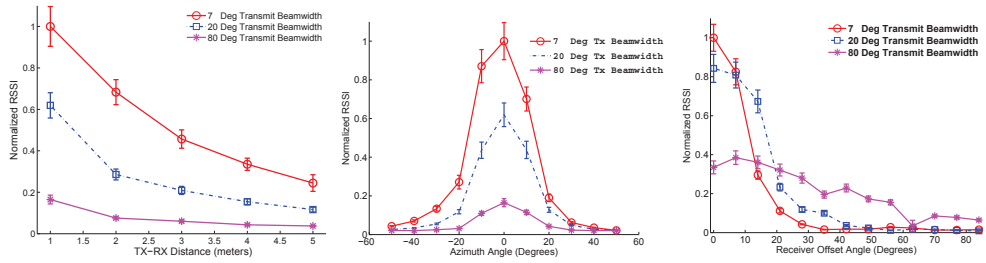

(a) RSSI versus inter-node distance (b) RSSI for various receiver orientations (c) Effect of receiver offset angle on RSSI

Figure 6: RSSI measurements for an unobstructed LOS 60 GHz link at different inter-node distances and relative angles

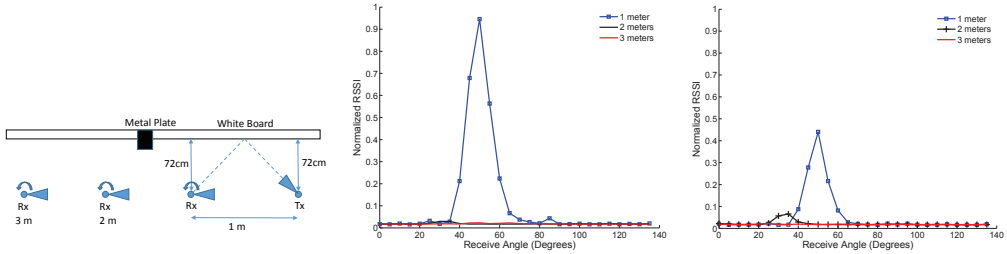

(a) Experimental setup (b) RSSI variation for different receiver orientations at 7° transmit beamwidth (c) RSSI variation for different receiver orientations at 20° transmit beamwidth

Figure 7: RSSI measurements for NLOS path resulting from reflection from a white board in a class room environment

The RSSI values are normalized to the average RSSI value for a LOS link at the same distance. We observe that the reflected path from the white board creates a strong link comparable to the LOS pathwith a smaller peak for the signal reflected from the metallic plate between the two boards. However, when the receiver is moved away from the transmitter, it is no longer aligned with the reflected path and there is no signal. Fig. 7c shows the results for 20° beamwidth. In this case, the reflected component from the board is significantly attenuated due to lower directivity gain. Moreover, there is no reflection component from the metal plate. However, due to larger angular spread of 20° beamwidth, the weak reflected component from the board is present at a distance of 2 meters as well. This experiment shows that strong reflected paths can originate from reflections from strong reflectors and these paths can be used to form alternate links upon blockage of LOS link. Moreover, the existence and strength of these paths depends on angular spread from Tx and Rx beamwidth and the angle of incidence.

4.3 WLAN Scenarios

Simulation Parameter	Value
Max. transmit slot	2ms
Beacon Interval	100 ms
Preamble Length	1.9 ns
Contention Slot	5 μs
SIFS	3 μs
DIFS	10 μs
Base Rate	27.5 Mbps
Highest Rate	6.7 Gbps
No. of MCS	13

Table 1: List of important simulation parameters

4.3.1 Setup

Here, we consider a wireless LAN scenario with nodes that are stationary or mobile, and capable of rotation and translation. For these experiments, we assume fully backlogged nodes with infinite data to send, and fixed data transmission slots (2 *ms* max TxOP) to send data at the selected rate. We model translation as random waypoint mobility, with random pause intervals (max. 1 sec) and human walking speed (1.5 m/sec). We use a similar mobility model for rotation, where the receiver picks a random angular distance to cover and a specified speed, after which it pauses for a random interval before rotating again. Translation and rotation are independent and their intervals can coincide. For channel models, we use the channel traces described above and the living-room channel model from 802.11ad evaluation methodology report [12] to explore fine grained steering, a larger set of Tx and Rx beamwidths, and rich propagation and reflection environments. In this model, obstacles like walls, windows and furniture lead to multiple reflected paths between communicating devices, in addition to the LOS path. Moreover, mobile obstructions in the environment, like humans, can lead to blocking of LOS or reflected paths. We use reference values for these reflection and penetration losses.

For performance comparison, we implement a baseline 802.11ad scheme, which uses SNR based receiver feedback for rate adaptation [3]. For successive losses which cannot be recovered by rate adaptation protocol, this scheme performs BFT with a fixed beamwidth. Hence it does not use Beam Sounding or beamwidth adaptation in comparison to MOCA. In later experiments, we also analyze the performance of this baseline scheme with beamwidth adaptation; however, the decision about rate or beamwidth adaptation is independent, in contrast to the *joint* adaptation in MOCA.

For these WLAN simulations, we use the PHY specifications in 802.11ad, including MCS schemes for OFDM and their corresponding sensitivity thresholds. Control packets are sent at the lowest bit rate (27.5 Mbps), and require minimum receive sensitivity (-78 dBm). OFDM achieves the highest data rate of 6.7 Gbps. We estimate the highest MCS supported on the link (and the highest rate) by using RSSI thresholds for each MCS and PHY layer, as given in the standard [4]. Packet header is considered to be received correctly if the remaining link budget is greater than -78 dBm (threshold for MCS-0, used to encode header) and payload is received correctly if link budget is greater than the threshold for MCS index used to encode data.

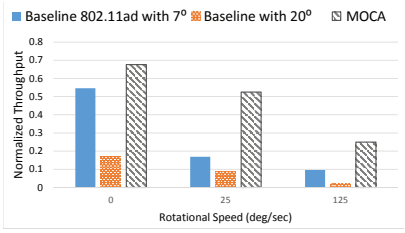

(a) Normalized throughput comparison

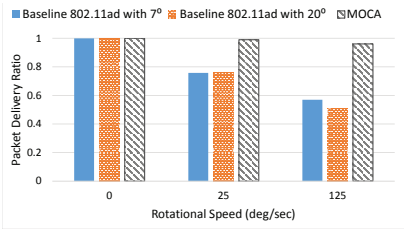

(b) Packet Delivery Ratio (PDR) comparison

Figure 8: Performance comparison of MOCA and Baseline 802.11ad schemes for trace-based WLAN emulation

Table 1 lists important simulation parameters.

4.3.2 *Trace-based Results*

We first analyze the performance of MOCA and the baseline scheme using the channel traces collected in the measurement study. Since we have limited beamwidths available via horn antennas, we consider the performance of the baseline scheme for both 7° and 20° beamwidths. Here we discuss the experiment for the lab environment, where we consider stationary nodes such that only the receiver is capable of rotation. This helps isolate the sector-misalignment issue and analyze the performance of various schemes. We consider the impact of nodal translation and environmental mobility in later experiments. Fig. 8a depicts the normalized throughput vs. receiver rotation speeds for the three schemes, whereas Fig. 8b shows the Packet Delivery Ratio (PDR) comparison for the same experiment. For each presented value, we perform Monte-Carlo simulations such that 95 percent confidence intervals are typically within 0.01 throughput and hence are not shown. We perform multiple simulations with random starting positions and present average results. In all our results, we normalize the throughput with the maximum possible throughput that is achieved with omniscient knowledge of the best sectors and highest rate on the link. Further, we consider a range of rotation speeds to understand how the performance changes as the extent of mobility increases.

For 0 deg/sec rotation, there is no nodal mobility and hence no misalignment; only the channel variations change the maximum MCS supported at the PHY layer. The baseline schemes achieve >95% PDR in this case, whereas MOCA gains slightly higher PDR due to fresh channel estimates in the Beam Sounding. However, the normalized throughput comparison for this no-mobility case shows differences among the three schemes due to MAC overhead, difference in directivity gains and MCS under-selection. The baseline with 20° achieves lowest throughput due to lower link budget, whereas the baseline with 7° achieves 82% of available throughput. MOCA slightly under-performs in this case, despite the higher PDR, due to increased overhead of the Beam Sounding. Observe that as the rotational speed increases, MOCA maintains >90% PDR for all rotation speeds. MOCA shows resilience to packet losses due to its Beam Sounding procedure, which probe the selected path prior to transmissions and identify link breakage in most cases. PDR for the baseline schemes degrades as the extent of mobility increases, due to higher losses to misalignment. Moreover, wider beamwidth provides greater resilience to misalignment, depicted by a higher PDR of the baseline 20° scheme compared to the baseline 7° scheme for moderate to high speeds. In terms of throughput, the baseline with 7° performs better than 20° at moderate speeds due to much higher achievable data rate. However, at higher speed of 360 deg/sec, its throughput becomes less than the baseline 20° due to higher frequency of misalignment and higher BFT overhead. Finally we observe that MOCA achieves higher throughput than the baseline schemes for all mobility scenarios, showing more than $2\times$ gains for the highest rotation speeds. Lastly, due to the large difference in link budget between 7° and 20° beamwidths and the improved PDR due to the Beam Sounding mechanism, MOCA selects 7° beamwidth throughout this experiment. Thus, to characterize joint adaptation, we next consider the 802.11ad evaluation channel model and use a larger codebook of available beamwidths.

4.3.3 *802.11ad Channel Model Results*

In the next experiment, we consider the 802.11ad living room channel model. We consider the receiver to be capable of rotation as well as translation. We consider the same rotation scenario as in the last experiment, while we model receiver translation as random waypoint mobility with human walking speed (1.3 m/sec). Rotation and translation motions are independent, and may overlap as well. For comparison to MOCA, apart from a baseline 802.11ad scheme with the narrowest beamwidth (3°, as specified in the standard), we also consider a scheme with a simple beamwidth adaptation strategy, which we term Widen Every Retransmission (WER), which exponentially backs-off in beamwidth selection to combat higher mobility; i.e., upon successive losses which cannot be recovered by the rate adaptation protocol, WER performs BFT with double the previous beamwidth. Furthermore, WER resets its beamwidth to the narrowest (3°) after every beacon interval.

Fig. 9 depicts performance comparison of the three schemes in terms of normalized throughput and PDR for a range of rotation speeds. We observe that all schemes obtain their highest throughput and PDR when the receiver is not rotating. Since the receiver still undergoes translation, resulting in misalignment and MCS changes, we expect gains for MOCA due to Beam Sounding and joint adaptation. However, MOCA only slightly outperforms the baseline schemes in this case. This limited MOCA gain is due to the additional overhead incurred by the Beam Sounding and slightly longer resilience training mechanisms.

Overhead Analysis: To investigate this overhead further, we compare the percentage overhead incurred by all three schemes in Fig. 10, which depicts the percentage of total transmission time spent by each scheme performing BFT (excluding the time for backoff and channel access) or Beam Sounding as well in case of MOCA. We plot the total overhead (Beam Sounding+Resilience Training) and Resilience Training overhead for MOCA on separate curves. For 0 deg/sec case, MOCA incurs higher training overhead compared to other schemes since it proactively searches for alternate fail-over paths. The overhead becomes even higher when we consider the additional time spent for Beam Sounding. However, this overhead is only 3.2% in this case, and is mitigated by the gains in terms of rate selection and link breakage identification, resulting in a better PDR and hence a slightly higher throughput for MOCA.

As the rotation speed increases, we observe that the training overhead for MOCA increases with a lower slope than the baseline schemes and is less than half as compared to WER scheme for moderate to high speeds. This reduction in overhead is partly due to MOCA's fast recovery mechanisms, thus reducing training

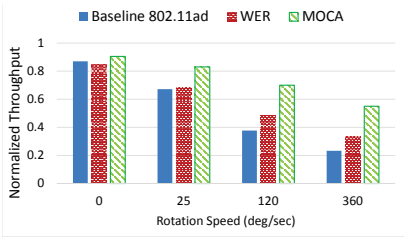

(a) Normalized throughput comparison

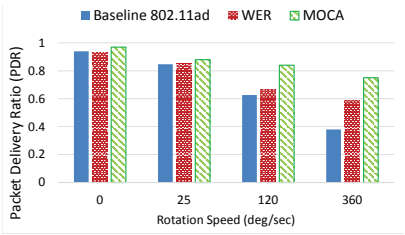

(b) Packet Delivery Ratio (PDR) comparison

Figure 9: Performance comparison of MOCA and Baseline 802.11ad schemes for living-room experiments

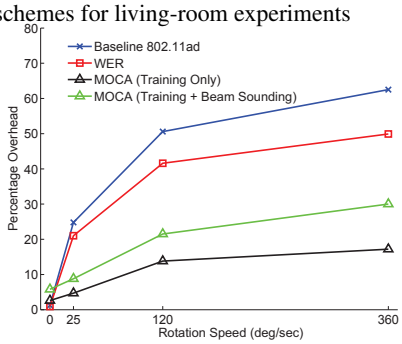

Figure 10: Overhead comparison for different schemes

frequency when successful, and partly due to its joint adaptation algorithm, which dynamically adjusts the beamwidth and rate to minimize this overhead.

The importance of beamwidth adaptation can further be understood by studying the normalized throughput and PDR at medium to high speeds in this experiment. WER shows higher resilience to misalignment and blockage, compared to the baseline scheme with narrowest beamwidth, by widening its beamwidth after detecting link breakage. This is reflected by its higher PDR and lower overhead compared to the baseline scheme. Although this simple beamwidth adaptation strategy improves performance, it is coarse grained and does not account for the current data rate or the extent of mobility in the network. MOCA shows further performance improvement compared to WER since its algorithm jointly adapts the beamwidth and rate based on recent history of packet losses to account for the extent of mobility. MOCA achieves highest throughput and PDR across all speeds, and shows more than $2\times$ gain compared to a fixed beamwidth 802.11ad scheme and up-to $1.5\times$ gain compared to a simple beamwidth adaptation strategy.

Blockage Experiments: In this experiment, we study the performance of MOCA under environmental mobility, which may lead to blockage. Measurement studies done in our hardware experiments as well as those in [12] and [21] suggest that strong reflected paths may result from reflections from shiny surfaces, metals, etc. In MOCA, antenna sectors are trained to use these paths as alternate links when the primary LOS link is blocked. In this case, the data rate on the alternate link depends on the nature of the reflecting surface and the angle of incidence. To study the impact of blockage

(a) Normalized throughput comparison

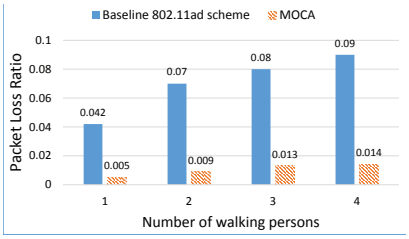

(b) Packet Loss Ratio comparison

Figure 11: Performance comparison of MOCA and baseline scheme under blockage from human mobility.

alone, we consider a point to point link between stationary nodes, separated by a distance of 7 meters. We use the reference model for living room environment from [12] with reflections from various surfaces, where obstacles with average human dimensions show random waypoint mobility at human walking speeds and random pause time.

Fig. 11a shows the comparison of normalized throughput and packet loss ratio between MOCA and the baseline 802.11ad scheme, as we vary the number of persons walking around in the room. We observe that a small number of packets are lost due to rate over-selection, but most occur due to breakage of links due to blockage, and the blockage induced packet loss gets higher as the number of walking persons increases. This is due to an increase in the blockage probability induced by human mobility. Packet loss ratio also increases for MOCA for an increase in human mobility; however, MOCA shows significant gains in terms of packet delivery ratio compared to the baseline scheme. MOCA achieves these gains by switching to alternate sectors upon blockage of the primary link, thus avoiding packet loss. In our experiments, we observe that the probability of alternate link being available is 0.82, 0.69, 0.61 and 0.55 for 1, 2, 3 and 4 walking persons, respectively. This decrease in alternate link availability results in MOCA throughput degradation as blockage probability increases. Fig. 11b shows the packet loss ratio for the two schemes. Although the throughput gains are not as pronounced as in misalignment experiments, due to lower frequency of blockage events compared to misalignment, MOCA shows significant improvement in packet delivery ratio compared to the baseline scheme.

5. RELATED WORK

Rate Adaptation: Many PHY rate adaptation protocols have been proposed, with the simplest using packet loss to adjust the MCS e.g., Auto Rate Fallback (ARF) [5]. Some protocols use explicit SNR or RSSI feedback from the receiver to adjust the data rate, e.g., Receiver Based Auto Rate [3] and the time-share-fair Opportunistic Auto Rate protocol [13]. A few protocols use PHY hints to adapt rate on a per-bit basis, e.g., SoftRate [20]. Recent work also addresses rate adaptation for multi-stream MIMO, e.g., [11] and [15]. However, in MIMO systems below 6 GHz, mul-

tiple streams and diversity are achieved exploiting the rich scattering propagation that is not present in the 60 GHz environment nor with switched narrow-beam antenna arrays. Thus, while any previous rate adaptation protocol can also be applied to 60 GHz WLANs, link degradation and breakage resulting from misalignment and blockage cannot be addressed by rate adaptation alone.

Directional and Beamforming Protocols: Prior work also considers blockage and directivity issues in 60 GHz WLANs. For example, [16] proposes a MAC protocol for 60 GHz networks which overcomes blockage by finding relay nodes to reach a blocked node. 802.11ad also allows the use of relay nodes, and [14] proposes extensions to this mechanism. While our approach employs reflected paths when the LOS path is blocked, relays can nonetheless be employed on top of MOCA when neither LOS nor reflected paths are available. Likewise, [17] studies the possibility of switching to widest sectors upon blockage detection to recover the link and concludes that beamwidth expansion can improve connectivity only in cases when the beam is partially blocked. In contrast, we employ beam widening in order to provide mobility resilience, including client rotation and translation. [18] proposes to switch to alternate fail-over sectors if the primary link fails, and develops a prediction framework for the quality of 60 GHz links. Such predictions can be used to refine the set of sectors that MOCA tests during MOCA's preemptive fast recovery phase. A sensor assisted beamforming protocol is presented in [22], where client's heading direction and speed, estimated by device sensors, is used by the AP to switch and adapt beams. This adaptation is based on serving the client at a minimum SNR. However, this approach does not help in recovering from blockage or finding non-LOS paths. Moreover, the sensor inputs are orders of magnitude slower than frame transmission times in 802.11ad and may not be effective in preventing packet loss. We infer mobility in MOCA based on Beam Sounding and frame feedback at packet level, and our beamwidth adaptation is based on maximizing average link throughput rather than instantaneous data rate. Lastly, prior work addresses the high overhead due to exhaustive search in BFT: [10] proposes out-of-band AOA estimation to reduce the search space in sector sweeps and [8] proposes an efficient beam switching technique to reduce BFT overhead. In contrast, our approach is to reduce the frequency of training by resilience and recovery mechanisms. Efficient BFT schemes can be used to further reduce the overhead in MOCA.

6. CONCLUSION

In this paper, we propose MOCA, a protocol for mobility resilience and fast link re-establishment in highly directional 60 GHz networks. We introduce Beam Sounding for quick identification of mobility-induced link breakage, and devise proactive recovery mechanisms by traversing codebook levels and fail-over sector search, with minimal overhead. MOCA is also the first protocol that targets maximizing average link throughput by switching antenna sectors and jointly adapting beamwidth and data rate depending on the extent of mobility in the network, while also incorporating the highly variable time required for control protocol overhead. Our experiments show that MOCA achieves up to $2\times$ throughput gains compared to a baseline WLAN scheme with in a diverse set of operational conditions.

7. ACKNOWLEDGEMENTS

The authors thank Joe Chen for his assistance in performing the experiments. This research was supported by Intel, the Keck Foundation, and by NSF grants CNS-1514285 and CNS-1444056.

8. REFERENCES

[1] K. Amiri, D. Shamsi, B. Aazhang, and J. R. Cavallaro. Adaptive Codebook for Beamforming in Limited Feedback MIMO Systems. In *Proc. of Annual Conference on Information Sciences and Systems (CISS)*, 2008.

[2] O. Bejarano, E. W. Knightly, and M. Park. IEEE 802.11ac: from channelization to multi-user MIMO. *IEEE Communications Magazine*, 51(10):84–90, 2013.

[3] G. Holland, N. Vaidya, and P. Bahl. A Rate-adaptive MAC Protocol for Multi-hop Wireless Networks. In *Proc. of ACM MobiCom*, 2001.

[4] IEEE 802.11 Working Group. IEEE 802.11ad, Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band. 2012.

[5] A. Kamerman and L. Monteban. WaveLAN-II: A High-performance Wireless LAN for the Unlicensed Band. *Bell Labs Technical Journal*, 2(3):118–133, 1997.

[6] A. Khattab, J. Camp, C. Hunter, P. Murphy, A. Sabharwal, and E. W. Knightly. WARP: A Flexible Platform for clean-slate Wireless Medium Access Protocol Design. *ACM SIGMOBILE Mobile Computing and Communications Review*, 12(1):56–58, 2008.

[7] H. H. Lee and Y. C. Ko. Low Complexity Codebook-Based Beamforming for MIMO-OFDM Systems in Millimeter-Wave WPAN. *IEEE Transactions on Wireless Communications*, 10(11):3607–3612, 2011.

[8] B. Li, Z. Zhou, W. Zou, X. Sun, and G. Du. On the Efficient Beam-Forming Training for 60GHz Wireless Personal Area Networks. *IEEE Transactions on Wireless Communications*, 12(2):504–515, 2013.

[9] T. Nitsche, C. Cordeiro, A. Flores, E. Knightly, E. Perahia, and J. Widmer. IEEE 802.11ad: Directional 60 GHz Communication for Multi-Gigabit-per-second Wi-Fi. *IEEE Communications Magazine*, 52(12):132–141, 2014.

[10] 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. of IEEE INFOCOM*, 2015.

[11] I. Pefkianakis, Y. Hu, S. H. Wong, H. Yang, and S. Lu. MIMO Rate Adaptation in 802.11n Wireless Networks. In *Proc. of ACM MobiCom*, 2010.

[12] E. Perahia. TGad Evaluation Methodology. http://www.ieee802.org/11/Reports/tgad_update.htm.

[13] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly. Opportunistic Media Access for Multirate Ad Hoc Networks. In *Proc. of ACM MobiCom*, 2002.

[14] S. K. Saha, L. Sun, and D. Koutsonikolas. Improving Connectivity, Coverage, and Capacity in 60 GHz Indoor WLANs Using Relays. In *Proc. of ACM Workshop on Wireless of the Students, by the Students, & for the Students*, 2015.

[15] W.-L. Shen, K. C.-J. Lin, S. Gollakota, and M.-S. Chen. Rate Adaptation for 802.11 Multiuser MIMO Networks. *IEEE Transactions on Mobile Computing*, 13(1):35–47, 2014.

[16] S. Singh, F. Ziliotto, U. Madhow, E. M. Belding, and M. Rodwell. Blockage and Directivity in 60 GHz Wireless Personal Area Networks. *IEEE Journal on Selected Areas in Communications*, 27(8):1400–1413, 2009.

[17] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan. Networking through Flexible Beams: A Link-Level Profiling. In *Proc. of ACM SIGMETRICS*, 2015.

[18] S. Sur, X. Zhang, P. Ramanathan, and R. Chandra. BeamSpy: enabling robust 60 GHz links under blockage. In *Proc. of USENIX NSDI*, 2016.

[19] Vubiq. V60WGD03 60 GHz Waveguide Development System. http://vubiq.com/v60wgd03/.

[20] M. Vutukuru, H. Balakrishnan, and K. Jamieson. Cross-layer Wireless Bit Rate Adaptation. *ACM SIGCOMM Computer Communication Review*, 39(4):3–14, 2009.

[21] H. Xu, V. Kukshya, and T. S. Rappaport. Spatial and Temporal Characteristics of 60-GHz Indoor Channels. *IEEE Journal on Selected Areas in Communications*, 20(3):620–630, 2002.

[22] Z. Yang, P. H. Pathak, Y. Zeng, and P. Mohapatra. Sensor-Assisted Codebook-Based Beamforming for Mobility Management in 60 GHz WLANs. In *Proc. of IEEE MASS*, 2015.