RICE UNIVERSITY

Overhead Constrained Joint Adaptation of MCS, Beamwidth and Antenna Sectors for 60 GHz WLANs with Mobile Clients

by

Muhammad Kumail Haider

A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE

Master of Science

APPROVED, THESIS COMMITTEE:

Dr. Edward W. Knightly, *Chair*
Professor of Electrical Engineering

Dr. Ashutosh Sabharwal
Professor of Electrical Engineering

Dr. Lin Zhong
Associate Professor of Electrical Engineering

HOUSTON, TEXAS
APRIL 2015

Abstract

Overhead Constrained Joint Adaptation of MCS, Beamwidth and Antenna Sectors for 60 GHz WLANs with Mobile Clients

by

Muhammad Kumail Haider

60 GHz directional networks pose new challenges in terms of rate selection and link maintenance in the presence of mobile nodes. In this thesis, we design and evaluate a novel cross-layer protocol, **BeamRAP**, for joint adaptation of antenna sectors, beamwidth and Modulation and Coding Scheme (MCS) for 60 GHz links with mobile clients. We propose this joint adaptation since beamwidth and alignment of directional antennas are the key determinants of link strength. This high directionality also introduces new challenges for link maintenance; e.g. link blockage and misalignment of directional antennas. Therefore, in BeamRAP, we introduce a new feedback mechanism to probe the link before transmissions, and devise an algorithm for joint adaptation of data rate and beamwidth. We also implement mechanisms to restore broken directional links with minimum overhead. We also build a 60 GHz programmable node and testbed using VubIQ 60 GHz transceivers, and conduct an extensive measurement study to collect channel traces over-the-air. Our experiments under multiple environmental and nodal mobility scenario show that BeamRAP achieves up to 2x gains in throughput as compared to a baseline 60 GHz scheme, which does not implement beamwidth adaptation.

Acknowledgements

In the name of Allah, the Beneficent, the most Merciful.

I extend my gratitude towards my thesis committee members, Dr. Edward Knightly, Dr. Ashutosh Sabharwal and Dr. Lin Zhong, for their guidance and help in this thesis work, as well as their valuable comments during the examination which helped me improve my thesis. I am especially thankful to my advisor, Prof. Edward Knightly for giving me the opportunity to do research at the Rice Networks Group, and helping me develop analytical skills which have been invaluable for me during the research work.

I am also grateful to my fellow researchers at RNG in particular and Rice in general, who gave me thoughtful insights about my work during discussions. They have truly enriched my learning experience at Rice.

I would also like to acknowledge Dr. Ihsan Qazi, my undergraduate advisor, who instilled this spirit of research and academic excellence in me. He has guided me even after I joined Rice, discussing key research problems and providing feedback about my current work, for which I will forever be indebted to him. I would also like to acknowledge my undergraduate school, Lahore University of Management Sciences, for providing me with an atmosphere conducive to research. LUMS is Life!

Last, but not the least, I would like to acknowledge my family, especially my parents, who truly are my role model in life. Belonging to a

small city in Pakistan, they supported me financially and always encouraged me to aim for the highest. I would never be at Rice if it was not for their relentless efforts and support.

Dedicated to my mother who played a vital role in inspiring my educational career!

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The 60 GHz frequency band, with its 7 GHz unlicensed spectrum, opens up avenues to multiGigabit communication. Since the wavelength in this band is on the order of a few millimeters, path loss is very high and directional antennas or electronically steerable beam-arrays are used to get directivity gain to attain the required link budget [26]. Therefore, 60 GHz links are highly directional, and require configuration of antennas at both the transmitter (Tx) and the receiver (Rx) for their establishment. For an electronically steerable array, beam patterns (or sectors) corresponding to different directions can be achieved by a codebook lookup of antenna weight vectors. A flexible codebook design also allows for beams with different beamwidths. As defined in IEEE 802.11ad standard for 60 GHz networks [5, 14], Beamforming Training (BFT) procedure is used to discover the direction (corresponding to virtual "sectors" that discretize beam steering) of the intended receiver. This process is described in detail in Fig. 1.1.

Beamforming Training helps in 60 GHz link establishment and selection of optimum sectors to achieve maximum data rate. However, environmental and nodal mobility may render this selection sub-optimal for subsequent transmissions. In particular, 60 GHz links are subject to two key mobility issues due to their directional

nature: *(i) Blockage* due to mobile obstructions (e.g. 15-20 dB path-loss in case of human body [23]); and *(ii) Misalignment* of the Tx and Rx sectors due to rotation or sideways mobility of nodes. In both cases, the data rate may significantly decrease or the link may completely break, and requires re-adjustment of sectors using BFT. However, BFT requires exhaustive search over all Tx and Rx sectors, and its timing overhead depends on sector beamwidths; the narrower the beamwidth, the longer the BFT time required. Therefore, there is a tradeoff between training a 60 GHz link for optimum direction and width in response to mobility vs. using the existing link (which may become sub-optimal) and avoiding the BFT overhead.

Due to the highly directional nature of 60 GHz links, the data rate depends on three factors: *(i)* the link budget associated with a Tx, Rx sector pair (e.g. the LOS paths result in links with higher budget than NLOS paths), *(ii)* the beamwidth used for Tx and Rx sectors, and *(iii)* the Modulation and Coding Scheme (MCS) used at the PHY. In the literature, several rate adaptation protocols have been proposed (details in Chapter 5) to adjust the MCS at the PHY in response to channel variations or nodal mobility. However, prior protocols are not designed to address how to adapt the latter two parameters i.e., the Tx, Rx sectors and their beamwidth. We propose that for maximizing throughput in highly-directional 60 GHz networks, it is necessary to select and adapt all three parameters jointly in response to channel and environmental variations.

In this thesis, we design, implement and evaluate a novel cross-layer protocol, **BeamRAP**, for overhead constrained joint adaptation of antenna sectors, beamwidth and MCS for long- term throughput maximization on 60 GHz links with mobile clients. Our design introduces a probing mechanism for 60 GHz links to adapt the antenna sectors and the MCS in response to mobility, minimizing the explicit BFT overhead. Moreover, BeamRAP is the first protocol to consider the problem of

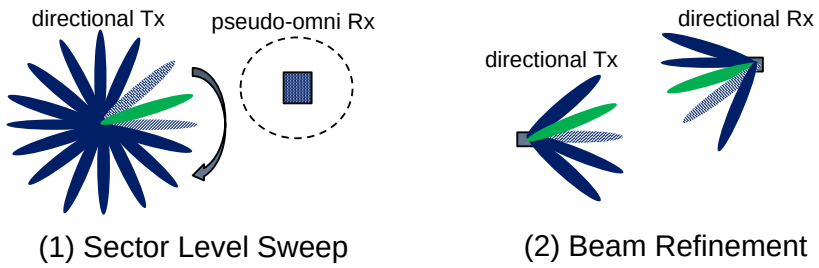

Figure 1.1: Two phases of **Beamforming Training** (BFT). 1) Sector Level Sweep (SLS) in which a node sweeps across all its Tx sectors, while the other receives in pseudo-omni mode and identifies the best Tx sector. 2) Beam Refinement Phase (BRP) for fine grained calibration of Tx and Rx sectors.

beamwidth and data rate adaptation *jointly*. The key idea is that the directivity in 60 GHz links introduces new determinants of the link strength (and thereby the data rate) compared to omni directional networks: the alignment of the Tx and Rx sectors and their beamwidth. A change in the link strength may be addressed by adapting the data rate, beamforming sectors and their beamwidth, or both. For example, if the distance between two nodes increases, the data rate may be reduced to overcome the reduced link strength. However, this problem can also be addressed by narrowing the beamwidth to achieve higher directivity gain while using the same rate. This modification of the data rate or beamwidth also impacts the long-term throughput of the link. Therefore, it is important to adapt data rate and beamwidth jointly, such that the throughput of the link is maximized.

This thesis is organised as follows. First, we describe the design of BeamRAP. In BeamRAP, we devise a new **Sector-pair Failure Inference** mechanism, which uses a probe-feedback exchange to identify blockage and misalignment on 60 GHz links. This exchange also helps in rate selection and adjustment of the Tx and Rx sectors prior to data transmissions by using short training and channel estimation fields, as described in Chapter 3. In case of link breakage, BeamRAP avoids any packet loss and uses the same transmit opportunity to attempt restoration of the link using its **Pre-emptive Fast Recovery** procedure. For this, we design a technique to identify

alternate fail-over sector pairs in advance, to be used in case of blockage, or to expand beamforming sectors to overcome misalignment.

We also devise an algorithm to select and adapt the beamwidth for the Tx and Rx sectors and the data rate, based on the channel estimates collected in probe feedback. This algorithm selects the optimum beamwidth and data rate to be used for the current data packet transmission, after the probe exchange takes place. This joint selection is necessary since there is a tradeoff between selecting narrower beams to get higher data rates vs. selecting wider beams for resilience against misalignment. The narrower the beamwidth, the higher the directivity gain and hence, the higher the data rates achievable. However, a narrower beamwidth also increases the number of sectors which have to be searched in BFT, thereby increasing the timing overhead. Moreover, narrower sectors are more susceptible to misalignment, as established in hardware experiments in Chapter 4. Therefore, narrower beamwidths can lead to significant overhead in BFT to overcome blockage and misalignment issues. BeamRAP estimates the rate of misalignment and blockage events using Sector-pair Failure Inference and then calculates the optimum beamwidth for each packet to maximize data rate and minimize BFT overhead.

Finally, we build a **60 GHz programmable node and 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 traces of channel strength and variations, as well as BER, for a lab and a classroom scenario. The VubIQ transceivers are equipped with horn antennas with 7°, 20° and 80° beamwidth and mechanical steering to emulate sector sweeps. Our measurements capture the impact of distance, misalignment, mobility and beamwidth on signal strength, and provide traces for evaluation of BeamRAP under realistic 60 GHz channel measurements. We also implement BeamRAP on a custom simulator

and evaluate its performance for various environments and mobility scenarios. We also implement a baseline 802.11ad scheme for comparison. Our experiments show that BeamRAP achieves up to 2x throughput gains compared to the baseline scheme in most of the blockage and misalignment scenarios.

Chapter 2

Background: IEEE 802.11ad

The IEEE 802.11ad standard defines modifications to the 802.11 Physical (PHY) and Medium Access Control (MAC) layers to enable operation in 60 GHz frequency band [5]. This chapter describes some of these specifications, which are relevant to BeamRAP design.

2.1 PHY Layer Specifications

IEEE 802.11ad defines four physical layers: Control PHY, Single Carrier PHY, OFDM PHY and Low Power Single Carrier PHY. Control PHY is used for all control and most management packets and because it uses the lowest bit rate (MCS-0) it requires minimum receive sensitivity (-78 dBm). All PHYs have a common preamble, which is composed of a short training field and a channel estimation field. It can be used for packet detection, synchronization, channel estimation and information about MCS used for payload.

As described in the introduction, directional antennas or electronically steerable antenna arrays are used in 60 GHz systems to achieve the required link budget. However, a pair of nodes must configure their transmit and receive antennas (or

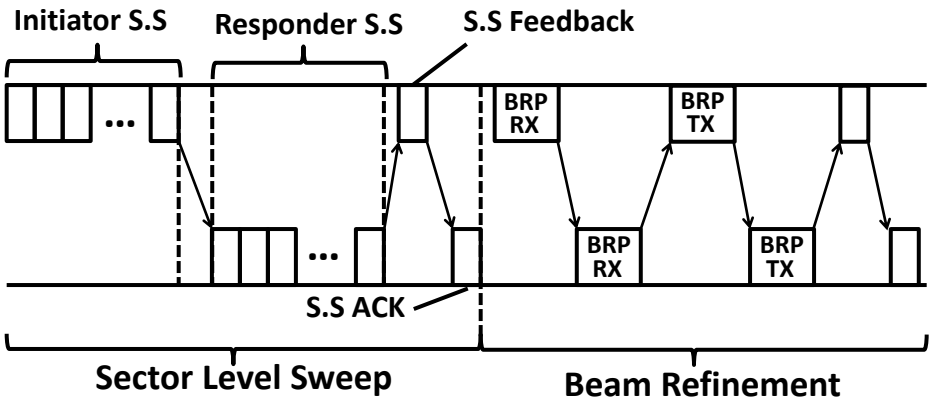

Figure 2.1: Beamforming Training (BFT) procedure

sectors) to participate in directional communication. Beamforming Training (BFT) is a bidirectional process in which the end nodes determine their optimum Tx and Rx sectors to reach each other by exchanging a sequence of training frames. As shown in Figure 1.1, BFT comprises of two phases: a mandatory Sector Level Sweep (SLS) phase, and an optional Beam Refinement Phase (BRP). During a sector sweep, the best transmit or receive sector is identified when the initiator switches across all available sectors while the responder receives in a quasi-omni pattern. However, this phase is coarse grained, and the refinement phase can be used to further divide the selected sectors to improve link budget.

2.2 MAC Layer Specifications

To support a variety of applications and traffic scenarios in 60 GHz systems, 802.11ad supports both random access and scheduled access. Channel time is divided into Beacon Intervals (BI), with a structure depicted in Figure 2.2. This structure specifies intervals for beacon transmission, beamforming, management frame exchanges and data transmission slots. Due to imperfect carrier sensing in directional networks,

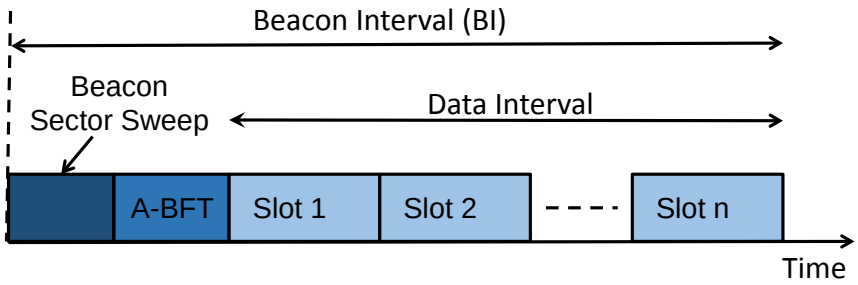

Figure 2.2: Phases of channel access in a Beacon Interval, as defined in 802.11ad. Beacon Interval starts with sector sweep of beacon frames by AP, A-BFT is used for beamforming during station association and is optional, followed by optional slot for management frame exchange. This is followed by data transmissions in Data Interval, which is further divided into (possibly multiple)periods of scheduled and random access.

802.11ad specifies a centralized control where the AP (in case of infrastructure networks) authenticates stations for any transmission. Moreover, in random access, nodes use directional transmission and reception to minimize interference and collisions. Therefore, in idle state, a node senses the channel in pseudo-omni mode until it receives an RTS or data packet. It can then switch to directional reception, using the sectors selected in BFT, to improve link budget. However, the transmission of control and data packets is always directional. For data transmission, the network can use scheduled access, contention based access, or both. In scheduled access, a data slot is pre-assigned to a station, whereas nodes can compete for contention based periods using modified 802.11 random access procedures. BeamRAP can be used for any of the access schemes specified. In any case, we introduce negative acknowledgements (NACK) in BeamRAP, as an enhancement to 802.11ad.

Chapter 3

BeamRAP Protocol Design

The key idea behind BeamRAP is that the highly directional nature of 60 GHz links introduces a new determinant of link strength: the beamwidth and alignment of directional antennas. The Tx and Rx antennas have to be perfectly aligned to achieve maximum data rates on the link. If this alignment is affected due to environmental or nodal mobility, the selected data rate is no longer supported on the link. Moreover, the width of the antenna sectors determines the directivity gain, and hence the maximum rates achievable on the link.

In this Chapter, we first analyze the design factors for BeamRAP and then describe the protocol's Joint Adaptation, Beam-pair Failure Inference, Pre-emptive Fast Recovery and Resilience Training mechanisms.

3.1 Design Factors

3.1.1 Link Failures in Directional Networks

The directional nature of links introduces new challenges in terms of link resilience and rate selection. Apart from channel degradation due to noise, interference and mobility

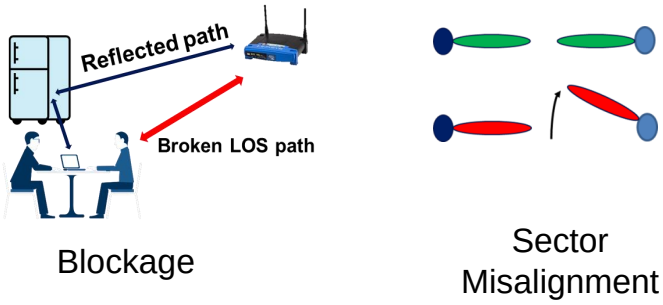

Figure 3.1: Packet losses due to directional nature of 60 GHz links. (a) *Blockage:* the existing link is obstructed by mobile objects, and alternate links have to be searched. And (b) *Sector Misalignment:* the Tx and Rx antenna sectors are misaligned and require re-adjustment.

etc., directional links are susceptible to breakage due to blockage and misalignment, as discussed below. It is important to identify and distinguish losses due to breakage from those due to unsupported MCS, since rate adaptation alone cannot recover broken links.

Collisions and interference can be made to be negligible in 60 GHz systems with sufficient directionality and centralized MAC design [6]. Therefore, we classify the major causes for 60 GHz channel degradation, which can change the optimal data rate or the optimal Tx and Rx sectors on a link, as follows:

- *Unsupported MCS:* The MCS supported at PHY layer can vary due to channel perturbations like noise and fading. It can also change due to nodal mobility; such that a receiver moving towards or away from the transmitter can respectively increase or decrease the maximum supported MCS.

- *Sector Misalignment:* Due to rotation of one of the nodes about its axis, or due to sideways mobility of one node relative to the other, the Tx and Rx sectors can become misaligned, thereby breaking the communication link.

- *Blockage:* If the LOS link is blocked by obstacles like humans (introducing 15-

20 dB loss [23]), it may also breaks the link and an alternate, non LOS link has to be explored. Blockage and sector misalignment scenario are illustrated in Figure 3.1. (improve the figure)

The unsupported MCS issue can be addressed using any of the existing rate adaptation protocols (some of them are discussed in Chapter 5); however, blockage and sector misalignment cannot be addressed by merely changing the MCS at the PHY layer. Therefore, the latter two scenarios require beam refinement to reconfigure the Tx and Rx sectors or complete beamforming training to find the alternate links respectively. One of the major challenges for BeamRAP design is to identify the cause of link degradation or link breakage, and then make a decision about whether to use rate adaptation or to align the antenna sectors again.

3.1.2 Time Bounded Beamforming Training

Narrower beamwidth at the Tx and Rx antennas helps achieve higher rates (by improving the link budget), but also introduces higher overhead for beamforming training.

If the beamwidth used at the Tx and Rx antennas is θ_{Tx} and θ_{Rx} respectively, the link budget (LB) is given as follows [7].

$$LB = \frac{40000}{\theta_{Tx}\phi_{vert}} + \frac{40000}{\theta_{Rx}\phi_{vert}} + P_{Tx} - PL(d) - \sum_{i=1}^{K} loss\ (obstacle\ i) \quad dBm \quad (3.1)$$

where ϕ_{vert} is the vertical beamwidth and is fixed. P_{Tx} is the transmit power, PL(d) is the path loss at distance d and K is the order of reflection. Therefore, the narrower the beamwidth used at the Tx and Rx antennas, the higher the link budget and the higher the data rates achieved. However, narrow beamwidth sectors require

more time for the exhaustive search in BFT in case the link breaks. Moreover, making the beamwidth narrower makes the link more susceptible to sector misalignment, as a slight rotation or sideways motion of the node can result in sector mis-match. More frequent misalignment instances lead to a higher frequency of Beamforming Training and hence a higher timing overhead.

As illustrated in Figure 1.1, Beamforming Training is required for link establishment and antenna re-adjustment so that nodes can select the best transmit and receive sectors requiring exhaustive search over all sector pairs in 802.11ad. Using 802.11ad timing values for its different slots, the length of a BFT slot in our implementation is as follows:

$$BFTtime = [\frac{a}{\theta_{Tx}} + \frac{a}{\theta_{Rx}} + c]\ \mu s \tag{3.2}$$

where a and c are protocol constants, with values 13320 and 145.31 μs in our implementation. Therefore, the time required for beamforming training varies inversely to the beamwidth, and for narrower beamwidths, it may become many times greater than the data transmission intervals. For example, for 3° beamwidth (minimum allowed in 802.11ad), the BFT requires 5 ms while the maximum transmit slot size is 2 ms. Moreover, in some cases, BFT can not be performed immediately after blockage since it requires authentication from a centralized controller (AP in infrastructure node) which may be delayed due to contention by other nodes in the network for the BFT. Initiating a sector sweep without scheduling increases interference and may collide with data transmissions.

For example, in 802.11ad, there are only two scenario when a node can initiate BFT: *(i)* At the beginning of a Beacon Interval, when the AP does its sector sweep for beacons, it may initiate BFT with a station or a station (or multiple stations) may contend to initiate BFT. *(ii)* In data transmission interval, a node may initiate

BFT instead of sending data, after it acquires the transmit slot through contention or scheduling. Therefore, non-availability of BFT slots or contention with other nodes may lead to a delay between a link breakage and the time BFT is initiated, resulting in an outage period where no transmissions are possible. In BeamRAP, we introduce a link recovery procedure to recover a broken link without requiring BFT, avoiding the training overhead and the outage period. It also reduces the interference footprint due to sector sweeps.

3.2 BeamRAP Protocol Design

In this section, we describe the design of BeamRAP protocol. The key idea in BeamRAP is to adapt the MCS, beamwidth and antenna sectors on 60 GHz directional links in response to mobility, such that the long-term link throughput is maximized. We introduce two new mechanisms in BeamRAP. First, we devise a Joint Adaptation metric to capture the rate vs. resilience tradeoff for joint selection and adaptation of beamwidth and data rate. This metric takes into account the maximum rate currently supported on the link, as well as the BFT overhead and the probability of link breakage associated with current beamwidth. We describe the protocol in detail in Section 3.2.2, outlining how these parameters are estimated. BeamRAP computes this metric prior to each data transmission for the entire set of available beamwidths, and selects the beamwidth and the corresponding rate which maximize this metric.

Second, we introduce a Pre-emptive Fast Recovery scheme to attempt the restoration of broken links (due to blockage or misalignment) in the same transmit opportunity. For this, we design a technique to identify alternate fail-over sector pairs in advance, to be used in case of blockage, or to expand beamforming sectors to overcome misalignment. If this pro-active scheme is successful, any outage period is avoided and communication is restored without explicit training overhead. We discuss this

procedure in detail in Section 3.2.3.

3.2.1 Joint-Adaptation Metric

Our metric for joint beamwidth and rate adaptation captures the tradeoff between selecting higher data rate due to directional gains of narrower beams vs. training and link re-establishment overhead due to mobility. Using renewal reward theory based analysis, the long-term average throughput (η), for a given beamwidth θ_i on a 60 GHz link is given by the following relation:

$$\eta(\theta_i) = \frac{(1 - P_{break})(Datarate)(TxTime)}{(1 - P_{break})(TxTime) + (P_{break})(BFTtime)} \tag{3.3}$$

$$. \qquad = \frac{[1 - P_{break}(\theta_i, \phi)][r(\theta_i)][t_{slot}]}{[1 - P_{break}(\theta_i, \phi)][t_{slot}] + [1 - P_{break}(\theta_i, \phi)][t_{BFT}(\theta_i)]} \tag{3.4}$$

where t_{slot} is the length of a data transmission slot, t_{BFT} is the time required for beamforming training for beamwidth θ_i, and we define P_{break} as the probability of a breakage event. A successful transmission results in a throughput gain at rate $r(\theta_i)$, whereas breakage leads to training penalty t_{BFT}, which also depends on beamwidth θ_i. If the value of these parameters is known, this metric can be used to maximize the long-term average throughput of the link by finding the beamwidth and the corresponding maximum rate which maximizes this metric. However, in a real network, the value of P_{break} is unknown and depends on the underlying environmental and nodal mobility.

For the estimation of maximum achievable rate ($r(\theta_i)$) and $P_{break}(\theta_i)$ in BeamRAP, we introduce a probe exchange mechanism prior to packet transmissions. As described in detail in Section 3.2.2, this probe exchange enables the estimation of link strength based on RSSI and also does refinement of beams to fine-tune the sectors. This

enables maximum rate selection for current beamwidth. Finally it helps to distinguish between losses due to unsupported MCS from those due to link breakage.

BeamRAP maintains a statistical estimate of $P_{break}(\theta_i)$ using a history based scheme; by counting the fraction of last n (n is a protocol parameter, which can be fixed or itself adapted in response to mobility) probes as being successfully transmitted, lost due to unsupported MCS or due to link breakage. This distinction is based on aforementioned probe-exchange based inference method. We then calculate the probability of breakage events ($P_{break}(\theta_i)$) as the fraction of last n probes being lost due to breakage ($P_{break} \in [0, 1]$ such that 0 means no breakage and 1 means link is broken for all probes).

Given the values of $P_{break}(\theta_i)$ and $r(\theta_i)$ for the current beamwidth θ_i, BeamRAP can estimate $\eta(\theta)$ for all $\theta_j \in [\theta_{min}, \theta_{max}]$ and switch beamwidth from θ_i to θ_j if $\eta(\theta_j) > \eta(\theta_i)$. For this estimation, we need to compute the rate and P_{break} for other beamwidths, given their values for the currently selected beamwidth θ_i. We calculate the data rate for θ_j from the current link budget (LB) for θ_i. If the beamwidth used for the Tx and Rx sectors is θ_{Tx} and θ_{Rx} respectively, the link budget (LB) on a 60 GHz link is given as follows [7].

$$LB = \frac{40000}{\theta_{Tx}\phi_{vert}} + \frac{40000}{\theta_{Rx}\phi_{vert}} + P_{Tx} - PL(d) - \sum_{i=1}^{K} loss\ (obstacle\ i) \quad dBm \quad (3.5)$$

For this, we can calculate the difference in directivity gain between θ_i and θ_j, add this difference to LB(θ_i) to calculate LB(θ_j), and then select maximum supported data rate at this link budget.

The value of P_{break} depends on the 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

by half, given the same rotational mobility conditions for the node. However, it also depends on the initial alignment of the beams. Since we perform beam refinement in a successful probe exchange, we assume perfect alignment of sectors for θ_i. To estimate values of $P_{break}(\theta_j)$ from $P_{break}(\theta_i)$, we use a first approximation assuming a uniform signal strength across the beam using the relation :

$$p_{SM}(\theta_j) = (\frac{\theta_i}{\theta_j}) * p_{SM}(\theta_i) \tag{3.6}$$

Based on our formulation, BeamRAP searches over the entire set of achievable beamwidths to select optimum beamwidth and rate. For computational efficiency, the search space can be restricted to neighboring beamwidths only to achieve incremental adaptation, and the computation can be made after specific breakage events.

3.2.2 Sector-pair Failure Inference

We introduce a probe based feedback mechanism in BeamRAP prior to data transmissions, such that a request and response is sent on the beamforming sectors explored and selected by the pair of nodes in the training process. We use a short structure for these probe packets, comprising only of a Short Training Field and a Channel Estimation Field apart from the PHY preamble. This probe exchange serves three important purposes in overall BeamRAP design; *(i)* it confirms the existence of a directional link corresponding to the selected sector pairs. If the probe exchange fails, BeamRAP initiates its Pre-emptive Fast Recovery procedure to restore the link. *(ii)* Short Training Field in the probe packets is used to perform beam refinement (specified in 802.11ad) prior to packet transmissions to fine tune the Tx and Rx sectors, without incurring overhead of explicit beamforming training. This ensures maximum link strength for the selected sectors. And *(iii)* Channel Estimation Field helps to estimate channel strength for rate selection.

Since probe request and response packets are encoded at base rate, they can be received successfully even if the selected rate for the data packets is unsupported on the link. Moreover, if the data is encoded at unsupported MCS, the receiver will fail to decode the data frame. However, the receiver can still *detect* the packet if the base rate is supported by the channel since the packet preamble is always encoded at MCS-0 with the lowest receive threshold. In this case, it sends a negative acknowledgement (NACK) to the transmitter. Moreover, in IEEE 802.11ad based systems, our probe feedback mechanism can be implemented by re-purposing RTS and CTS control packets with the aforementioned two fields. However, these packets are now used for probing the link and not for channel reservation. These probe packets are 10 μs long, much shorter than BFT and independent of beamwidth.

As shown in Algorithm 1, BeamRAP initiates a transmit opportunity by sending a probe packet using the primary Tx sector, discovered earlier in the Resilience Training. In case of a successful exchange of probe packets, BeamRAP obtains a fresh estimate of the channel $RSSI_{curr}$ and selects rate based on the estimate $RSSI_{est}$ for data transmission. If a packet is *detected* but not decoded, the receiver can still measure the RSSI and reply with a NACK. BeamRAP uses this feedback in the NACK to differentiate two scenarios. If $|RSSI_{curr} - RSSI_{est}| < 3 * \sigma_{RSSI}$(where σ_{RSSI} is standard deviation in RSSI), BeamRAP infers loss due to higher MCS. However, if $|RSSI_{curr} - RSSI_{est}| \geq 3 * \sigma_{RSSI}$, BeamRAP infers partial blockage and calculates rate and beamwidth for next transmission using its joint adaptation algorithm, as described below. A partial blockage may result from obstacles with lower penetration loss so that the link is not completely broken, or due to slight misalignment such that the link is not broken but the supported rate degrades significantly. Selecting rate based on the most recent value in this case allows the beamwidth adaptation mechanism (described in next section) to initiate BFT, if optimum. In case of no probe

response, BeamRAP declares link breakage and initiates its fast recovery mechanism.

3.2.3 Pre-Emptive Fast Recovery

In case of link breakage resulting from blockage or sector misalignment, 60 GHz links require training of beamforming sectors to re-establish connection. However, beamforming training is a bidirectional process and requires authentication from the AP, as well as coordination between the two nodes. Non-availability of BFT slots or contention with other nodes may lead to a delay between a link breakage and the time the training is initiated, resulting in an outage period where no transmissions are possible. In BeamRAP, we introduce a link recovery procedure to recover a broken link without requiring BFT, avoiding the training overhead and the outage period, when successful.

The key idea is to pro-actively search for alternate links in advance, which can be used later in case of blockage. If blockage results from environmental mobility, there is a probability that the alternate links still exist and can be used opportunistically. We propose **Resilience Training** procedure to search for these alternate links by modifying the beamforming training procedure. The first phase in Resilience Training is also a sector sweep, in which the initiator exhaustively sweeps across its Tx sectors i.e. virtual sectors created by predetermined antenna weight vectors (AWVs), while the responder is in pseudo-omni reception. However, in BeamRAP, the responder identifies two strongest sectors and this information is fed back to the initiator. This is followed by a sector sweep by the responder. In the second phase, the two nodes repeatedly refine their AWVs for not only the strongest sector pairs, but the second best sector pairs as well. The former is selected as primary link for data transmissions, whereas the latter is cached as an alternate link to be used upon failure of the primary link.

Therefore, if the probe exchange on the primary sector pairs fails, BeamRAP attempts to recover the link from misalignment or blockage by expanding the beamforming sectors or by switching to alternate links. These proactive measures, if successful, avoid the wastage of current transmit opportunity and any outage period when the nodes try to discover each other again. For the protocol, if the probe request on the primary link fails, BeamRAP attempts at a fast recovery of the link by first expanding the beamforming sectors. The transmitter also sets a flag in the probe request indicating the receiver to expand its sectors as well. If a response is received, the nodes switch to this expanded sector as their primary link and select rate based on the most recent RSSI.

If sector expansion fails to restore the link, BeamRAP probes the alternate link discovered in the training to recover from breakage. If the alternate link is available, the receiver can identify that probe request corresponds to the alternate link, based on the mutually agreed sector ID in BFT. It can then send a response on the alternate link as well. If this response is received successfully, the nodes switch to this alternate link and resume data communication. It also gives a hint that link breakage happened due to blockage. In case the alternate link also fails, BeamRAP requests to schedule BFT at MAC layer. These steps are shown in detail in Algorithm 1.

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 BeamRAP procedures. Note that there is a possibility that the alternate link is also blocked by the obstruction blocking the primary link, or that the alternate link is lost due to environmental or nodal mobility. In this case, recovery attempt using the alternate link will not succeed. However, since more than 90% of Beamforming Training comprises of sector sweep [5], which is the same as in 802.11ad

BFT, the timing overhead of this proactive alternate link search is low. We evaluate this overhead, the probability of alternate link being available, and its recovery gains in Chapter 4.

Algorithm 1 : Beam-pair Failure Inference

$loss\_index = 0$ *% number of packets lost*
$Probe\_fail = 0$ *% counter for failed probe packets*
$BFT\_flag = 0$ *% flag to indicate BFT required*

$\rightarrow$ *Send Probe Request using primary Tx sector*
if (*Response is received*) **then**
  $RSSI_{est} = \alpha * RSSI_{est} + (1 - \alpha) * RSSI_{curr}$
  $MCS = MCS\_lookup(RSSI_{est})$ *% select rate*
  $\rightarrow$ *send Data*
  if (*ACK is received*) **then**
    $\rightarrow$ *end transmit opportunity*
  else if (*NACK is received*) **then**
    if $(|RSSI_{curr} - RSSI_{est}| < 3 * \sigma_{RSSI})$ **then**
      % declare loss due to higher MCS
      $RSSI_{est} = \alpha * RSSI_{est} + (1 - \alpha) * RSSI_{curr}$
    else if $(|RSSI_{curr} - RSSI_{est}| \geq 3 * \sigma_{RSSI})$ **then**
      % declare loss due to partial blockage
      $RSSI_{est} = RSSI_{curr}$
    end if
  else
    $loss\_index++$
  end if
else
  % declare loss due to blockage or misaligned sectors
  $loss\_index++$
  $Probe\_fail++$
  $backoff$
end if

Algorithm 2 : Pre-emptive Fast Recovery

continued from Algorithm 1

if $(RTS\_fail >= 2)$ **then**

Sector Expansion

CW= rand(0,31) % double the backoff window

– wait for $DIFS + CW * slot\_time$

– double the Tx sector width and resend RTS requesting receiver to double its Tx sector width for CTS as well

if $(CTS\ is\ received)$ **then**

– send Data

Identify previous loss due to misalignment

% set double Tx sector width as default

% reset params: $RSSI_{est}$, σ_{RSSI} *and* MCS_{guard}

else

Alternate Link Probing

CW = rand(0,63) % double the backoff window

– wait for $DIFS + CW * slot\_time$

– use alternate Tx sector explored in BFT to resend RTS, request receiver to use alternate sector too[1]

if $(CTS\ is\ received)$ **then**

– send Data

Identify previous loss due to blockage

% set alternate Tx sector as default

% reset params: $RSSI_{est}$, σ_{RSSI} *and* MCS_{guard}

else

% declare link recovery failure

$BFT\_flag = 1$ *% BFT required*

$loss\_index++$

end if

end if

end if

[1]Since Tx,Rx pairs are explored and agreed upon at both ends, sender can tell the receiver which sector to use for CTS.

Chapter 4

Implementation and Evaluation

In this Chapter, we describe our 60 GHz measurement testbed, and implementation and evaluation of BeamRAP.

4.1 Implementation and Evaluation Methodology

To evaluate BeamRAP on a realistic 60 GHz channel, and to capture effects of LOS path, misalignment of antenna beams and reflections on signal strength, we design and implement a 60 GHz programmable node and testbed using VubIQ 60 GHz transceivers (compliant with 802.11ad) [9]. Our testbed also implements mechanical beam-steering and uses multiple horn antennas to achieve directivity gains at different beamwidths. Using this testbed, we conduct an extensive measurement study to collect over-the-air traces of channel strength and variations for a lab and a classroom scenario.

4.1.1 FPGA Platform for Over-the-Air Measurements

For our 60 GHz channel measurement study, we use a mm-Wave development platform from VubIQ [9]. The platform consists of a transmitter and a receiver waveguide

system operating in 57-64 GHz unlicensed frequency band and has 1.8 GHz modulation bandwidth. We generate I/Q baseband signal at different modulations and rates using Wireless Open-Access Research Platform (WARP) [10]. WARP is a custom built FPGA platform, including the MAX2829 chipset that provides RSSI readings. We use WARPLab [12], 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 [11] using a 6 GHz Ultra Dynamic Range Differential Amplifer (ADL5565). On the receiver side, the signal goes through the VubIQ receiver module, a subtractor circuit and the WARP board. Using WARPLab, we dump over-the-air channel measurements into a buffer, which can later be analyzed. To achieve different sector widths, we use horn antennas with 7°, 20°, and 80° beamwidths at both the transmitter and the receiver. We also use an omni-directional antenna to achieve pseudo-omni reception. Moreover, to emulate sector sweep as described in IEEE 802.11ad, we achieve mechanical beam-steering by mounting the transceivers on a programmable rotating table which can emulate fixed sectors. This setup is also used to emulate misalignment and nodal rotation scenario. Our measurement testbed is shown in Figure 4.1.

4.1.2 Trace Driven WLAN Simulation Platform

We further implement BeamRAP in a custom simulator and use the results of the over-the-air measurement study to drive trace-based simulations. Moreover, to explore a broader set of operational conditions including electronic beam-steering, a larger set of beamwidths, and multiple environmental scenarios, we implement the channel model used in the 802.11ad channel evaluation methodology report [8]. This model covers path loss, channel variations, antenna gains, reflection and penetration losses in a

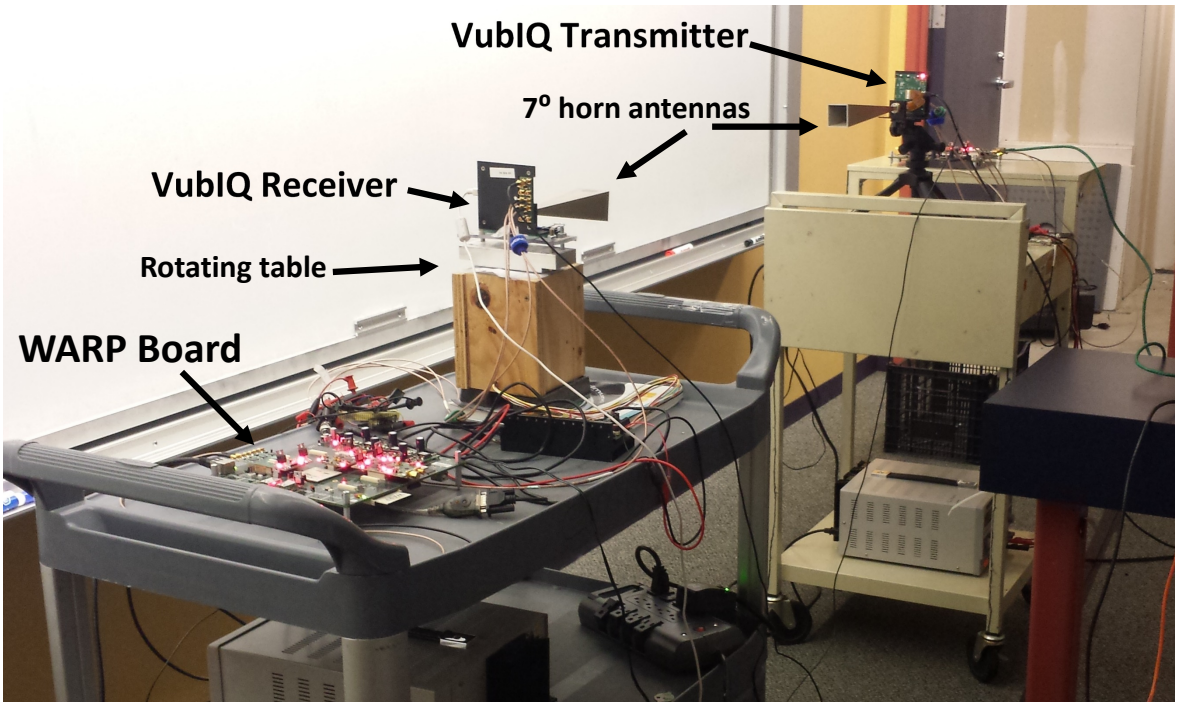

Figure 4.1: 60 GHz channel measurement experimental setup, with two VubIQ transceivers connected to WARP platforms. The transceivers are equipped with horn antennas of varying beamwidths and can be mechanically steered.

living-room and office environment. The experiments characterize the impact of nodal mobility (translation and rotation), environmental mobility, sector-misalignment and human blockage.

The simulator uses ray tracing and link budget analysis to model signal propagation and detection. For our trace-based simulations, we use the same link budget values as measured in our hardware experiments. Similarly, we model variation in RSSI on the channel as a random variable with same mean and variance as those from the measurements. 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 our measurement data set). Since our measurement data points are relatively coarse-

grained (compared to continuous time mobility and rotation in the simulator), we use weighted average to calculate RSSI values at these intermediate points. Moreover, in the trace based simulations, we are limited to 7,20 and 80 degree beamwidths, as mentioned earlier.

4.2 Over-the-Air Measurement Study

4.2.1 LOS Path and Receiver Rotation

In this experiment, we measure the signal strength and channel variations for a point-to-point link with an unobstructed LOS path. The experiments are performed in an electronics laboratory environment with dimensions 8 x 4 meters. The transmitter position 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. For LOS measurements, we move the receiver away from the transmitter, along a straight line to a distance of 5 meters. We take measurements at intervals of 1 meter, such that we have set of 5 measurement positions. Moreover, to capture the impact of sector-misalignment, we rotate the receiver along the azimuth, from $-170°$ to $170°$ (with $0°$ pointing in transmitter direction), and take measurements at intervals of 10 degrees. Each measurement consists of 100 packet transmissions over-the-air, and we record RSSI values for each run. We repeat the same experiment for transmitter beamwidths of $7°$, $20°$, and $80°$ to study differences in directivity gain, whereas the receiver beamwidth is fixed at $20°$.

Figure 4.2a shows the variation of RSSI as the inter-node distance increases from 1-5 meters, for all three transmit beamwidths. We have normalized the RSSI values with respect to the maximum ($7°$ beamwidth at 1 meter distance) for comparison. $7°$ beamwidth achieves the maximum signal strength, due to its higher directivity

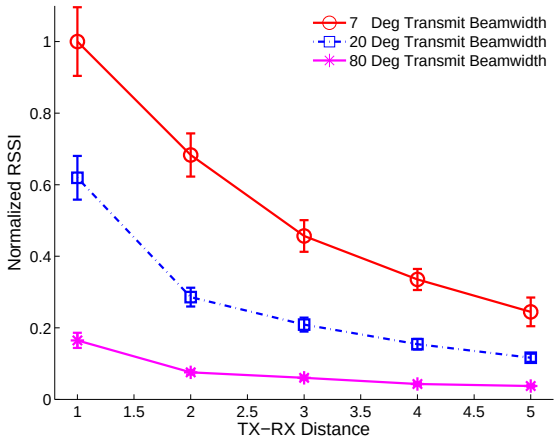

(a) Variation of RSSI w.r.t. inter-node distance

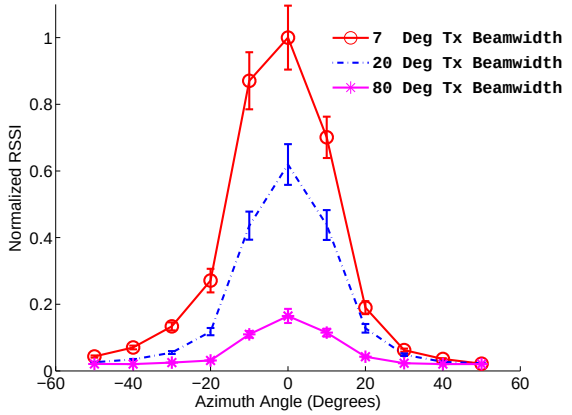

(b) RSSI variation for different receiver orientations

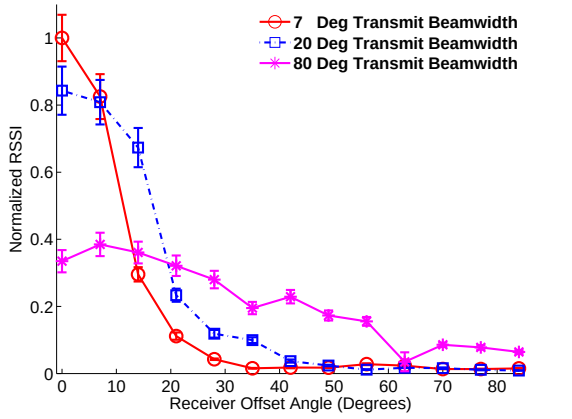

(c) Effect of receiver offset angle on RSSI

Figure 4.2: RSSI measurements for an unobstructed LOS 60 GHz link at different inter-node distances and relative angles

gain, across all distances. At 1 meter, the difference in signal strength between 7° and 20° is 12.2 dB, whereas 20° beamwidth has 11.71 dB higher signal strength than 80°. This difference in directivity gains is within 3 dB of the theoretical gains given by Equation (1). Figure 4.2b shows the received power versus receiver antenna angle for for all three transmit beamwidths, while keeping receiver beamwidth at 20°, at inter-node distance of 1 meter. Here the RSSI values are shown for an angular spread of −50° to 50°, since there is no signal reception beyond this range. The results illustrate the dominance of the LOS path and presence of weaker NLOS components, especially for 7° beamwidth. Moreover, the angular spread of 7° is the largest compared to wider beamwidth sectors. The reason is that we rotate the receiver with fixed beamwidth, while 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.

4.2.2 Radial Mobility and Misalignment

To analyze the effect of radial mobility and the resulting misalignment between transmit and receive antennas, we move the receiver along the circumference of a quad circle, with the transmitter placed at the center of the circle. The receiver beamwidth is fixed at 7° in this experiment, whereas the transmitter assumes all three available beamwidths in successive experiments. At 0° angle on the circle, the transmitter and the receiver are perfectly aligned. We then move the receiver along the circumference, and take measurements at every 7° separations, so that receiver is placed at 1 meter radial distance, at angles 0°, 7°, 14° and so on. The results for this experiment are depicted in Figure 4.2c. In this experiment, the misalignment results from radial mobility of receiver around the transmitter. Therefore, we observe that signal

strength 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 shows that while narrower beamwidths provide higher directivity gains and data rates, are much more susceptible to misalignment and hence link breakage due to radial or rotational mobility. Wider beams provide maximum resilience and uniform signal strength across larger spread of Tx-Rx relative angles.

4.3 Reflection Experiments

We conduct experiments for the measurement of non Line of Sight paths resulting from reflection off of various objects. Here we present an experiment in a classroom environment, where we measure the RSSI on an NLOS link resulting from reflection off a white board. The experimental setup is shown in Figure 4.3a, where the transmitter is pointing diagonally towards the white board, at a distance of 72 cm from the board. The receiver is placed at a distance of 1 meter from the transmitter, in a straight line, and 72 cm from the board. 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. We then move the receiver along a straight line, away from the transmitter, and take similar measurements at intervals of 1 meter. We perform separate experiments for 7° and 20° transmitter beamwidths, while the receiver beamwidth is fixed at 7 degrees.

Figure 4.3b shows the variation in received signal strength vs. receiver angles between 0° and 135° for transmitter beamwidth of 7 degrees, at different distances. The RSSI values are normalized to the average RSSI value for a LOS link at same

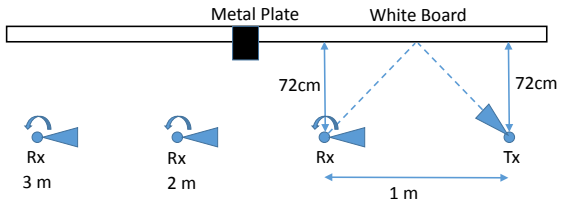

(a) Experimental setup

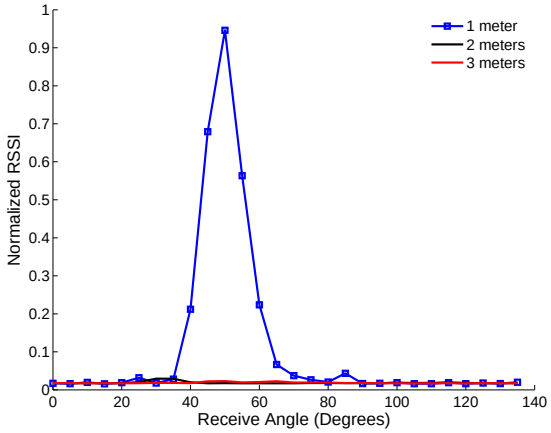

(b) RSSI variation for different receiver orientations at 7° transmit beamwidth

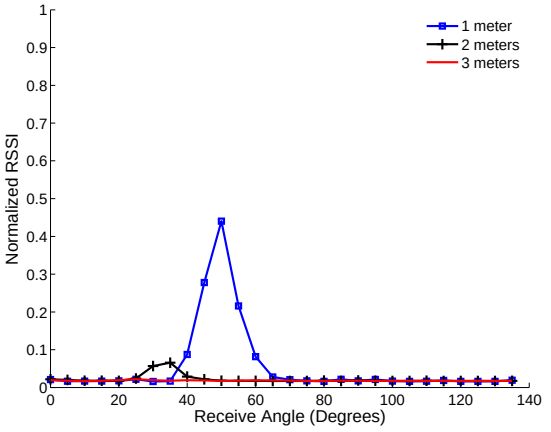

(c) RSSI variation for different receiver orientations at 20° transmit beamwidth

Figure 4.3: RSSI measurements for NLOS path resulting from reflection from a white board in a class room environment

distance. We observe that the NLOS path from reflection from the white board creates a very strong link, with strength comparable to the LOS path. We also observe 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 we observe no signal. Figure 4.3c shows the results for the same experiment with 20° beamwidth. We observe that the NLOS component from the board is much attenuated due to lower directivity gain. Moreover, there is no reflection component from the metal plat. However, due to larger angular spread of 20° beamwidth, we observe that a weak NLOS component from the board is present at a distance of 2 meters as well. This experiment shows that strong NLOS paths can originate from reflections from strong reflectors and these paths can used to form alternate links upon blockage of LOS link. We also observe that the existence and strength of these paths depends on angular spread from transmit and receive beamwidth and angle of incidence.

4.4 WLAN Results

4.4.1 Setup

Here, we use the aforementioned data sets as well as 802.11ad working group 60 GHz channel models to explore throughput in WLAN scenarios incorporating both MAC- and PHY-layer dynamics. For performance comparison, we implement a baseline 802.11ad scheme. The baseline scheme uses Auto Rate Fallback for rate adaptation [13]. In case of excessive losses (four successive attempts at reducing MCS, or twelve packets), this scheme falls back to beamforming training to recover from blockage or misalignment. Moreover, this scheme does not have beamwidth adaptation and beamwidth remains the same throughout the experiment. For evaluation, we consider

normalized throughput and packet delivery ratio (PDR) as performance metrics. We define normalized throughput as the maximum possible throughput that is achieved when simulator has omniscient knowledge of the best sectors to use, as well as the highest rate achievable on the resulting link. We consider both point-to-point links and WLANs with 10 stations.

We also implement a living-room model based on 802.11ad evaluation methodology [8] to explore fine grained steering, a larger set of transmit and receive beamwidths, and dfferent propagation and reflection environments. The link budget for this theoretical model is same as given in eq. 1. We use reference values from measurement study in [8] to quantize loss in signal strength due to reflection and penetration, depending on the nature of object and the angle of incidence. We use link-budget analysis to calculate signal strength, where the initial link budget of a ray is the sum of transmit power and directivity gains from transmit and receive beamforming. The narrower the transmit and receive sectors, the greater the beamforming gain. As the signal propagates, it incurs reduction in link budget due to path loss, channel variations, penetration and reflections off walls and obstacles. For simplicity, we only consider first and second order reflections. The relation to calculate directivity gain as a function of horizontal beamwidth ${\Theta_{HP}}^{\circ}$ and vertical beamwidth ϕ_{HP}°, as given in [7], is:

$$D = \frac{4000}{{\Theta_{HP}}^{\circ}\phi_{HP}^{\circ}} \tag{4.1}$$

and for path loss:

$$PL(d)[dB] = PL(d_{\circ}) + 10nlog_{10}(d) \tag{4.2}$$

where $PL(d_{\circ})$ is 68 dB and 80 dB for LOS and NLOS respectively and we assume $\phi_{HP}^{\circ} = 60^{\circ}$ for simulations.

Packet reception is determined based on remaining link budget and angle of arrival, when a ray hits the receiver. We implement the same PHY specifications as given in 802.11ad in Chapter 2, including MCS schemes for OFDM and their corresponding sensitivity thresholds. 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 [5]. Packet header is considered to be received correctly if 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 threshold for MCS index used to encode data.

Our simulator models stationary objects, as well as mobile obstacles such as humans (penetration loss of 20 dB), which can show translational, rotational or radial mobility. The reflection and penetration characteristics of obstacles depend on their nature and angle of incidence of the signal (reference values taken from [8]). We also use *Random Waypoint Mobility* to model obstacle movement, receiver translation and rotation.

Table 4.1 gives a list of important simulation parameters.

Simulation Parameter	**Value**
Packet length	262143 Bytes
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 4.1: List of important simulation parameters

4.4.2 Nodal Mobility Experiments

60 GHz links are susceptible to link breakage due to misalignment of Tx and Rx sectors resulting from nodal mobility. Moreover, translation can also result in change in link strength, apart from channel variations, which affects the data rate and packet reception. In these experiments, we study the performance of BeamRAP and baseline 802.11ad scheme in nodal mobility scenario. The transmitter is fixed, whereas the receiver is mobile and can show translation as well as rotational motion. We model both these motions as Random Waypoint Mobility with fixed speed (human walking speed i.e. 1.31 m/sec and different rotation speeds), and random pause time for translation. We present results for two different propagation environments.

In the first experiment, we consider a lab scenario based on our measurement traces and realistic 60 GHz channel measurements. Figure 4.4 shows the normalized throughput vs. receiver rotation speeds for three schemes; baseline 802.11ad with 7° transmitter beamwidth, baseline 802.11ad with 20° transmitter beamwidth and BeamRAP. The receiver beamwidth is fixed at 20°. We observe that 20° beamwidth leads to very low signal strength and thereby, lower rates. Therefore, its normalized throughput across all rotation speeds is very low. When the receiver does not rotate, we observe that BeamRAP achieves approximately 68% of available throughput. The loss in throughput is partly due to MAC overhead and under-selection by the rate adaptation, and dominantly due to BFT overhead and loss in directivity gain upon sector expansion. Note that although the receiver does not rotate, its translation can also result in sectors being misaligned, though less frequently. We observe that the baseline scheme with 7° beamwidth also achieves 55% of maximum achievable throughput. BeamRAP drives its gains over the baseline scheme due to sector expansion mechanism, which can avoid an outage period when the link breaks due to misalignment. This is also reflected in the packet delivery ratio comparison in Fig-

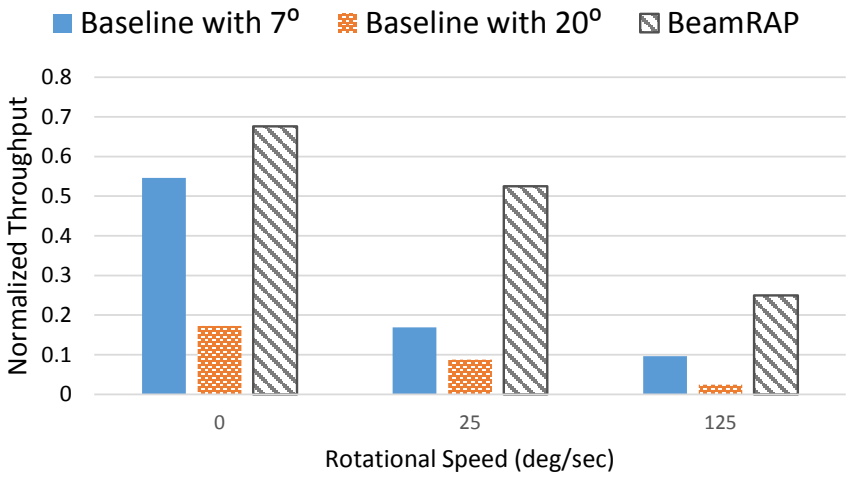

Figure 4.4: Normalized throughput comparison for BeamRAP vs. baseline scheme with different beamwidths, across a range of receiver rotation speeds. BeamRAP achieves more than 2x gains over baseline scheme at moderate speeds.

ure 4.5. We observe that BeamRAP achieves 100% PDR at moderate speeds. At higher rotation speeds, baseline scheme suffers throughput degradation and packet losses due to higher misalignment frequency and outage periods when contending for BFT. BeamRAP's throughput degrades with speed at a lower rate, such that it achieves more than double the throughput of baseline scheme at higher rotation speeds. We also observe in our experiments that due to huge difference in channel strength between 7° Tx beamwidth and 20° beamwidth, BeamRAP selects 7° beamwidth throughout the experiment.

For the second experiment, we consider the living room scenario presented in 802.11ad evaluation methodology report. To capture the impact of receiver rotation alone, we fix the position of the receiver as well, such that it rotates about its axis at variable speed and random pause time. The beamwidth for the baseline scheme is fixed at the narrowest (3°). Moreover, we also implement a version of BeamRAP without beamwidth adaptation, such that its beamwidth is also fixed.

Figure 4.6 shows the comparison of the three schemes for different rotational speeds. We observe that when the receiver is not rotating (no mobility), baseline

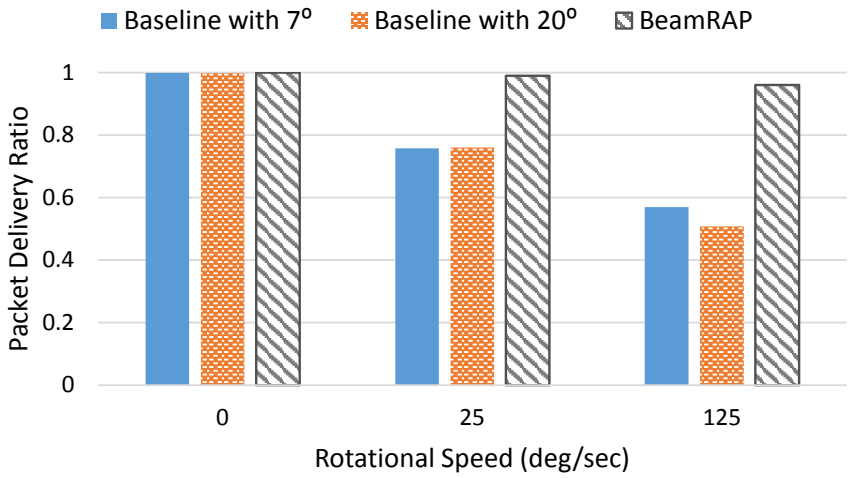

Figure 4.5: Packet Delivery Ratio comparison for BeamRAP vs. baseline scheme with different beamwidths, across a range of receiver rotation speeds. BeamRAP achieves 100% PDR for moderate speeds. Baseline scheme suffers higher loss ratio due to misalignment and outage periods.

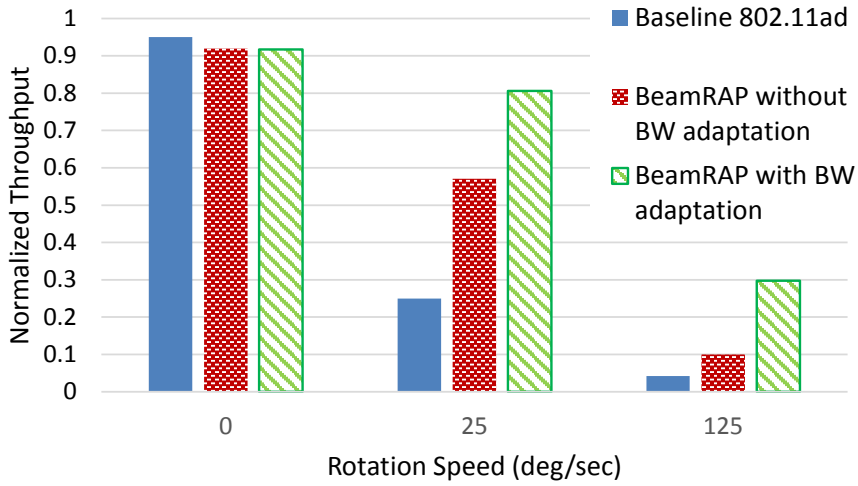

Figure 4.6: Normalized throughput comparison for BeamRAP (with and without beamwidth adaptation) vs. baseline scheme, across a range of receiver rotation speeds. BeamRAP achieves more than 2x gains over baseline scheme at moderate speeds.

scheme performs the best, achieving 95% of available throughput. BeamRAP achieves approximately 91%, slightly less than baseline scheme due to additional overhead of extended BFT procedure and MAC overhead. For a receiver rotation speed of 25 deg/sec, the performance of baseline scheme drops significantly due to sector misalign-

ment issues. The normalized throughput is around 0.25 due to frequent outages and timing overhead of BFT. We observe significant performance gains in both BeamRAP schemes. Without beamwidth adaptation, BeamRAP achieves double the throughput of baseline scheme due to identification of misalignment events and sector expansion. This experiment also shows the significant gains achieved when beamwidth is adapted jointly with rate. The difference between three schemes is even greater in case of high rotation speed 125 deg/sec (nearly the speed for quickly bringing the phone from pocket to ears).

4.4.3 Blockage Experiments

In this experiment, we study the performance of BeamRAP and the baseline scheme under environmental mobility, which may lead to blockage of the directional link. Measurement studies done in our hardware experiments as well as those in [8] and [1] suggest that strong NLOS paths may result from reflections from shiny surfaces and metals etc. The antenna sectors can be 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 angle of incidence. To study the impact of blockage alone, we consider a point to point link between stationary nodes, separated by a distance of 7 meters. We then analyze throughput and packet loss ratio due to blockage by simulating human blockage. We use the reference model for living room environment from [8] with reflections from various surfaces, where obstacles with average human dimensions show random waypoint mobility at human walking speeds and random pause time.

The normalized throughput for BeamRAP and the baseline scheme is shown in Figure 4.7, as we vary the number of persons walking across the room. The higher the number, the greater the blockage probability and the longer the blockage duration. In

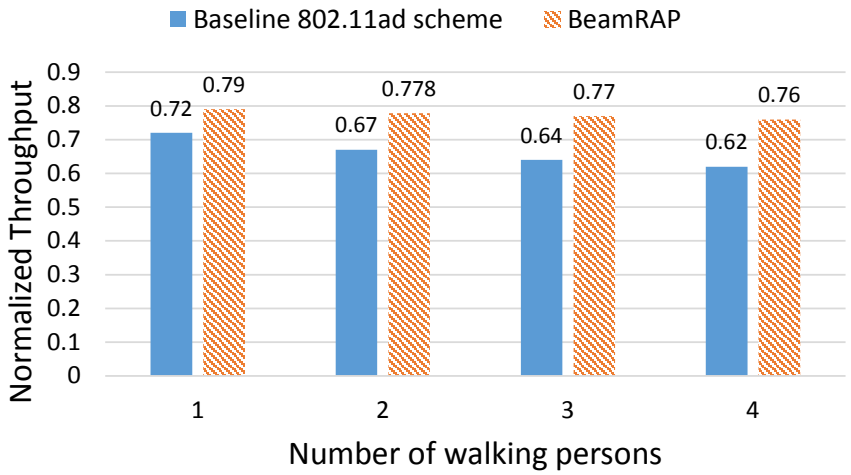

Figure 4.7: Normalized throughput comparison for BeamRAP vs. baseline scheme under blockage from human mobility.

this experiment, we observe that BeamRAP achieves higher throughput compared to the baseline scheme, and the gains increase as the blockage probability and duration increases. BeamRAP achieves these gains by switching to alternate sectors upon blockage of the primary link, thus avoiding BFT overhead. However, we observe that as the number of obstacles increases, there is a higher probability that the alternate link is also blocked by the same or some other obstacle, in which case, BeamRAP also falls back to requesting BFT. In our experiments, we observe that the probability of alternate link being available is 0.82, 0.69, 0.61 and 0.55 for 2, 4, 6 and 8 persons walking in the room, respectively. This decrease in alternate link availability results in BeamRAP throughput degradation as blockage probability increases. Figure 4.8 shows the packet loss ratio for the two schemes. Although the throughput gains are not as pronounced as in misalignment experiments, due to lower data rates on alternate links, BeamRAP shows significant improvement in packet delivery ratio under blockage, compared to the baseline scheme.

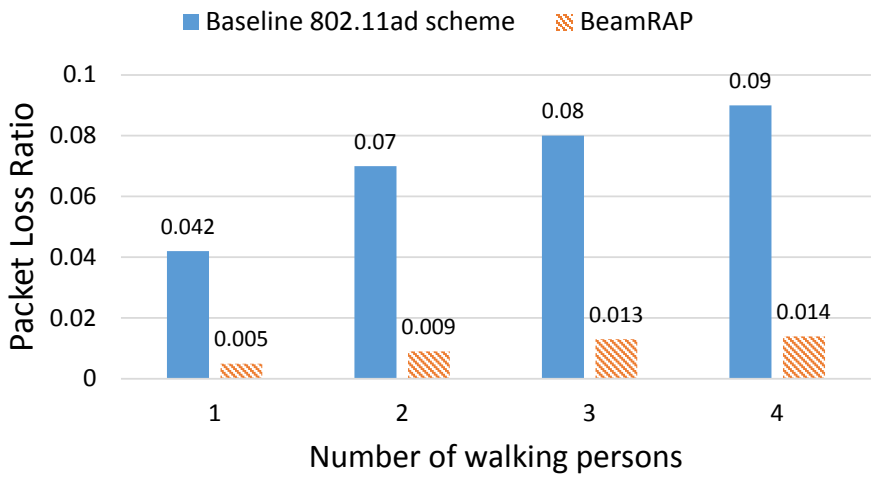

Figure 4.8: Packet loss ratio comparison for BeamRAP vs. baseline scheme under blockage from human mobility.

Chapter 5

Related Work

5.1 Rate Adaptation Protocols

Previous work in MCS adaptation can be classified into three categories: *(i) Loss Triggered protocols* use reception of ACKs at the link layer to modify the MCS, e.g., Auto Rate Fallback (ARF)[13] and adaptive ARF [15]. *(ii) SNR Triggered Protocols* use explicit SNR or RSSI feedback from the receiver to adjust the data rate, e.g., Receiver Based Auto Rate [16]. *(ii) Cross Layer protocols* use PHY hints and incremental redundancy to adapt rate on a per-bit basis e.g., SoftRate [17] and Strider [19]. These protocols have their pros and cons in terms of required feedback and accuracy of loss to rate mapping, as studied in [20]. However, these protocols are designed and evaluated for omni-directional networks, and do not address issues resulting from directivity of directional links. Some protocols address rate adaptation in multi-stream MIMO, e.g., [21] and [22]. However, in MIMO, multiple streams are achieved using spatial diversity and not due to narrow-beam directional antennas. Since RSSI feedback is required in 802.11ad and has less computation overhead than SNR, BeamRAP uses it to adapt data rate. It can use any of aforementioned protocols for adapting MCS as well. However, it differs in terms of selecting Tx, Rx sectors

and their beamwidth as well, which is not addressed by existing protocols.

5.2 Directional and Beamforming Protocols

Prior work addresses blockage and directivity issues in directional networks. For example, work in [23] develops a MAC protocol for 60 GHz networks which overcomes blockage by finding relay nodes to reach a blocked node. The AP uses sector sweep to find direction of all stations, and the stations also gather this information and feed back to the AP. Kim et al. discuss the problem of rate adaptation for directional multicast in 802.11ad networks [24]. However, they do not design any MAC or beamforming protocol, and assume that the AP has the knowledge of all the stations in each sector. [25] discusses how to adapt the beamwidth of an AP to maximize channel utilization and satisfy the required link budget criterion. However, it does not address beamwidth adaptation between two nodes for data transmissions.

Thus, BeamRAP is the first protocol to address rate selection in 60 GHz networks by using rate adaptation and beamforming training jointly, such that it adapts the Tx, Rx sectors, their beamwidth and MCS.

Chapter 6

Conclusion

We have designed a novel protocol (BeamRAP) for joint beamforming and rate adaptation in 60 GHz systems. Since directional links in 60 GHz band are highly susceptible to blockage, BeamRAP proactively searches for alternate links via a new training mechanism. In case of packet loss, it uses variance in RSSI feedback and probe RTS packets to differentiate between losses due to higher MCS encoding, blockage and sector misalignment. Its Pre-emptive Fast Recovery method uses sector expansion and alternate links explored in extended BFT to restore connection with the receiver. Moreover, it dynamically adjusts the beamwidth of Tx and Rx antennas based on the frequency of blockage and misalignment events, to minimize the Beamforming Training overhead. Our evaluation of BeamRAP shows that under link deterioration, blockage and sector misalignment scenarios, it shows up to 2x gains as compared to a baseline scheme which does not use any beamwidth adaptation. These results reiterate the need to do joint beamforming and rate adaptation in 60 GHz networks.

Bibliography

[1] Xu, H., Kukshya, V., Rappaport, T. S. , "Spatial and temporal characteristics of 60-GHz indoor channels," Selected Areas in Communications, IEEE Journal on, 20(3), 620-630. 4.4.3

[2] Yang, L. Lily, "60GHz: opportunity for gigabit WPAN and WLAN convergence," ACM SIGCOMM Computer Communication Review, 2008.

[3] Rappaport, Theodore S, "Wireless communications: principles and practice," Vol. 2. New Jersey: Prentice Hall PTR, 1996.

[4] J. Kim, S. Kim, S. Choi, and D. Qiao, "CARA: Collision-aware rate adaptation for IEEE 802.11 WLANs," in Proc. IEEE INFOCOM, 2006.

[5] IEEE 802.11 working group, "IEEE 802.11ad, Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band," December 2012. 1, 2, 3.2.3, 4.4.1

[6] C. Cordeiro, D. Akhmetov, and M. Park, "IEEE 802.11ad: introduction and performance evaluation of the first Multi-Gbps WiFi technology," in Proc. 2010 ACM Int. Workshop mmWave Commun.. 3.1.1

[7] Yong, S. K., Xia, P., Valdes-Garcia, A, "60GHz Technology for Gbps WLAN and WPAN: from Theory to Practice," 2011. 3.1.2, 3.2.1, 4.4.1

[8] Perahia, E., "TGad Evaluation Methodology," IEEE-09/296r8, July 2009 4.1.2, 4.4.1, 4.4.1, 4.4.3

[9] Vubiq. V60WGD03 60 GHz Waveguide Development System. [Online]. Available: http://vubiq.com/v60wgd03/ 4.1, 4.1.1

[10] Rice University WARP project. Available at:http://warp.rice.edu. 4.1.1

[11] Analog Devices ADL5565-EVALZ. 4.1.1

[12] N. Anand, E. Aryafar, and E. W. Knightly, "WARPlab: a fexible framework for rapid physical layer design. In Proc. of ACM workshop on Wireless of the students, by the students, for the students," 2010. 4.1.1

[13] Kamerman, Ad, and L. Monteban, "WaveLANII: a high performance wireless LAN for the unlicensed band," Bell Labs tech. journal 2.3 (1997) 4.4.1, 5.1

[14] T. Nitsche, C. Cordeiro, A. Flores, E. Knightly, E. Perahia and J. Widmer, "IEEE 802.11ad: Directional 60 GHz Communication for Multi-Gbps Wi-Fi," IEEE Communications Magazine, 52(12):132-141, December 2014. 1

[15] M. Lacage, M.H. Manshaei, and T. Turletti, "IEEE 802.11 Rate Adaptation: A Practical Approach," in Proc. ACM MSWiM 2004. 5.1

[16] G. Holland, N. Vaidya, and P. Bahl, "A Rate-Adaptive MAC Protocol for Multihop Wireless Networks," in Proc. ACM MobiCom 2001. 5.1

[17] M. Vutukuru and H. Balakrishnan, and K. Jamieson, "Cross-Layer Wireless Bit Rate Adaptation," in Proc. ACM SigComm 2009. 5.1

[18] J. Camp and E. Knightly, "Modulation Rate Adaptation in Urban and Vehicular Environments: Cross-Layer Implementation and Experimental Evaluation," Mobicom 2008.

[19] Gudipati, Aditya, and Sachin Katti, "Strider: automatic rate adaptation and collision handling," ACM SIGCOMM 2011. 5.1

[20] Biaz, Saad, and Shaoen Wu, "Rate adaptation algorithms for IEEE 802.11 networks: A survey and comparison," Computers and Communications, 2008. 5.1

[21] I. Pefkianakis, Y. Hu, S. H. Wong, H. Yang, and S. Lu, "MIMO Rate Adaptation in 802.11n Wireless Networks," in MOBICOM, 2010. 5.1

[22] W.-L. Shen, Y.-C. Tung, K.-C. Lee, K. C.-J. Lin, S. Gollakota, D. Katabi, and M.-S. Chen, "Rate Adaptation for 802.11 Multiuser MIMO Networks," in Proc. of ACM MobiCom, 2012. 5.1

[23] S. Singh, F. Ziliotto, U. Madhow, E. Belding, and M. Rodwell, "Blockage and directivity in 60 GHz wireless personal area networks: from cross-layer model to multihop MAC design," IEEE J. Sel. Areas Commun., vol. 27, no. 8, Oct. 2009. 1, 3.1.1, 5.2

[24] W. Kim, T. Song, and S. Pack, "Rate adaptation for directional multicast in ieee 802.11 ad networks," in Consumer Electronics (ICCE), 2012 IEEE International Conference on. IEEE, 2012, pp. 364365. 5.2

[25] K. Chandra, R. Prasad, I. Niemegeers, and A. Biswas, "Adaptive beamwidth selection for contention based access periods in millimeter wave wlans," in Consumer Communications and Networking Conference (CCNC), 2014 IEEE 11th, Jan 2014, pp. 458464. 5.2

[26] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson et al., "Demystifying 60GHz Outdoor Picocells," in Proc. ACM MobiCom, 2014. 1

[27] Ardalan Amiri Sani, Hasan Dumanli, Lin Zhong, and Ashutosh Sabharwal, "Power-efficient directional wireless communication on small form-factor mobile devices," in Proc. ACM/IEEE Int. Symp. Low Power Electronics and Design (ISLPED), August 2010.