A Detailed Characterization of 60 GHz Wi-Fi (IEEE 802.11ad)

by

Daniel Holmes-Mitra

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Electrical and Computer Engineering

Waterloo, Ontario, Canada, 2017

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

The emergence of wireless local area network (WLAN) standards and the global system of mobile communication (GSM) in the early 1990s incited tremendous growth in the demand for wireless connectivity. Iterative technological enhancements to cellular and WLAN improved wireless capacity and created a breadth of new mobile applications. The continued increase in display resolutions and image quality combined with streaming displacing satellite/cable has created unprecedented demands on wireless infrastructure. Data-caps on cellular networks deter over consumption and increasingly shift the growing burden to Wi-Fi networks. The traditional 2.4/5 GHz Wi-Fi bands have become overloaded and the increasing number of wireless devices in the home, public, and workplace create difficult challenges to deliver quality service to large numbers of client stations. In dense urban areas, the wireless medium is subjected to increased interference due to overlapping networks and other devices communicating in the same frequency bands. Improvements to conventional Wi-Fi are approaching their theoretical limits and higher order enhancements require idealized conditions which are seldom attainable in practice.

In an effort to supplant to scaling capacity requirements a very high frequency WLAN amendment has been proposed (IEEE 802.11ad). IEEE 802.11ad, also referred to as Wireless Gigabit (WiGig), operates in the globally unlicensed 60 GHz band and offers channel bandwidths nearly 100x as wide as 802.11n. The higher bandwidth facilitates multi-Gbps throughput even with the use of lower complexity modulation coding schemes (MCS). IEEE 802.11ad relies heavily on rate adaptation and high beamforming gain to mitigate interference and fading as signals in the 60 GHz band suffer from higher atmospheric absorption and free space path loss (FSPL). Due to the unique nature of 60 GHz wireless there have been numerous research efforts. Many studies have been directed at simulation and modeling of the 60 GHz channel. However modeling the channel is difficult as real-world environments are highly dynamic with varying link quality and conditions which cannot be accurately predicted by conventional techniques. Some research is focused on medium access control (MAC) enhancements to improve overall capacity by coordinating concurrent links or reducing communication overhead for example. Lastly, there has been a limited amount of real world testing of 802.11ad due to lack of availability of commercial platforms and measurement instrumentation. Some researchers tested early generation devices in certain use cases such as in vehicles for media streaming, in data centers to augment the wired network, or in basic indoor and outdoor environments.

This research contains two main components. In the first study, analytical models are applied to estimate line of sight (LOS) 802.11ad performance for realistic antenna parameters. The second part contains a comprehensive evaluation of performance and reliability

of early generation 802.11ad hardware. This characterization emphasizes environmental performance (e.g. conference room, cubical farm, open office), multiple-client testing (multiclient), multiple network interference (spatial re-use), and stability in the presence of station mobility, physical obstructions, and antenna misalignment. In order to evaluate 802.11ad, early generation platforms from technology vendors were used in extensive test suites. The hardware tested included docks for wireless personal area networking (WPAN) applications, client laptop stations, and reference design access points (APs). Finally, a customized proof-of-concept (PoC) platform was engineered which allowed finer control over front end antenna configuration parameters such as: topology, placement and orientation. The PoC also served as a suitable means to identify practical limitations and system design engineering challenges associated with supporting directional multi-Gbps (DMG) communication in the 60 GHz band.

Acknowledgements

I would like to thank Cisco Systems® for sponsoring my research and allowing me to complete two summer internships wherein I had access to early generation 802.11ad devices and extensive testing resources and facilities.

I would also like to thank my managers Stephen Saliga and John Blosco for their guidance and mentorship throughout my internships. Without their motivation and support none of this would have been possible. When I met them I was novice in RF theory and Wi-Fi but with your patience and guidance I gained a deep understanding of a very fascinating domain.

I would like to thank Waterloo University and my supervisor Dr. Gordon Agnew for his support and for helping me to gain access to leading edge technologies from the private sector.

Finally, I would like to thank my girlfriend and family for supporting me through this long journey of higher education.

Dedication

This thesis is dedicated to those who tinker.

Table of Contents

List of Tables

List of Figures

List of Acronyms

A-BFT Association Beamform Training

AC Access Category

ACK Acknowledgement

AES Advanced Encryption Standard

AP Access Point

AR Augmented Reality

ATI Announcement Transmission Interval

BFT Beamform Training

BHI Beacon Header Interval

BI Beacon Interval

BSS Basic Service Set

BTI Beacon Transmission Interval

CBAP Contention Based Access Period

CBC Cipher Block Chaining

CCA Clear Channel Assessment

CCM CBC Counter Mode

CEF Channel Estimation Field

COTS Commercial Off The Shelf

CPHY Control PHY

CSMA/CD Carrier Sense Multiple Access with Collision Detection

CTS Clear to Send

CW Contention Window

DC Data Center

DMG Directional Multi-Gigabit Per Second

DTI Data Tramission Interval

EDCA Enhanced Distributed Channel Access

EIRP Effective Intensity of Radiated Power

EM Electromagnetic

ESS Extended Service Set

FCC Federal Communications Commission

FCS Frame Check Sequence

FST Fast Session Transfer

GCM Galois-Counter Mode

GSM Global System for Mobile Communication

HCF Hybrid Coordination Function

IEEE Institute of Electronics and Electrical Engineers

IFS Inter-frame Spacing

IoT Internet of Things

IP Internet Protocol

ITU International Telecommunications Union

ISM Industrial Scientific and Medical

LPSCPHY Low Power Single Carrier PHY

LOS Line of Sight

MAC Medium Access Control

MIMO Multiple Input Multiple Output

MPDU MAC Protocol Data Unit

MSDU MAC Service Data Unit

MTU Maximum Transmission Unit

MU Multiple User

NAV Network Allocation Vector

NLOS Non Line of Sight

NSS Network Subsystem

OFDM Orthogonal Frequency Division Multiplexing

OTA Over the Air

PAL Protocol Adaptation Layer

PHY Physical Layer

PLCP Physical Layer Convergence Protocol

PoC Proof of Concept

PPDU MAC Protocol Data Unit

PSDU MAC Service Data Unit

QAM Quadrature Amplitude Modulation

QoS Quality of Service

RF Radio Frequency

RRM Radio Resource Management

RSSI Received Signal Strength Indicator

RTS Request to Send

SCPHY Single Carrier PHY

SP Service Period

SQI Signal Quality Indicator

SSID Service Set Identifier

SSW Sector Sweep

SSW-FB Sector Sweep Feedback

STA Station

STF Short Training Field

SU Single User

TCP Transmission Control Protocol

TDMA Time Division Multiple Access

ToR Top of Rack

TP Throughput

TXOP Transmit Opportunity

UDP User Datagram Protocol

VHT Very High Throughput

VR Virtual Reality

V2X Vehicle to Device and Vehicle to Infrastructure

WBE Wireless Bus Extension

Wi-Di Wireless Display

WiGig Wireless Gigabit - Synonymous with 802.11ad

WLAN Wireless Local Area Network

WPAN Wireless Personal Area Network

YOY Year over Year

Chapter 1

Introduction

In 1990 the Institute of Electronics and Electrical Engineers (IEEE) founded working group 802.11 to draft standards for Wireless Local Area Networks (WLAN). Around the same time the Global System for Mobile Communication (GSM) was defined and the first GSM phones were approved for sale in 1992. In 1993, code-division multiple-access (CDMA) spread-spectrum digital cellular systems were deployed in the US[1]. The evolution of digital wireless communication was set in motion and sparked a new age of technologies such as mobile telephony in the form of voice calling and text messaging. In 1997, the legacy 802.11 WLAN protocol was released supporting data rates of 1-2 Mbps in the 2.4 GHz Industrial Scientific and Medical (ISM) band, primarily for use in industrial business applications. Over the next two decades, the 802.11 standard was amended to increase wireless capacity, reliability and range which incited a breadth of new applications. 802.11b was introduced in 1999 offering data-rates up to 11 Mbps which became highly popular for consumer applications. Concurrently the Wi-Fi Alliance formed as a trade association to hold the Wi-Fi trademark used for all Wi-Fi ® products today. The Wi-Fi Alliance tests Wi-Fi equipment for IEEE standard compliance to ensure interoperability, backward compatibility and promote WLAN[2]. Despite such a brief evolution Wi-Fi enabled devices are now so ubiquitous that its difficult to imagine a time without the wireless connectivity known today. The continued evolution of both WLAN and cellular standards is paramount to satisfy the inevitable and insatiable growth in wireless demand.

1.1 Motivation

The Growth of Wireless Load

Wi-Fi is a prolific technology and the demand for wireless connectivity is growing at unprecedented rates. By 2016 the number of installed Wi-Fi devices matched the human population with a total of 12 billion Wi-Fi certified devices shipped, and this is expected to surpass 15 billion before 2017[3]. By 2024 this is expected to grow to 27 billion connected devices; a testament to the prevalence of wireless applications[4]. Wi-Fi usage is dominant in all sorts of portable electronics such as smartphones, tablets and laptops. Smartphones consume 92% of global mobile data traffic, despite only making up 18% of handsets in use globally, many of which are low throughput (TP) GSM. Furthermore, Cisco predicts an average 46% growth year over year (YOY) for tablets; the fastest growing device adoption with data growth more than doubling by 113% annually[5]. Modern surveys show that mobile data consumption is growing 50% YOY and data used over Wi-Fi networks accounts for over half of the total mobile data consumption and is steadily growing[6].

In the beginning wireless demand required much less bandwidth as it consisted primarily of text based email and messaging services. Today mobile electronics have become a technological pillar for all ages and are frequently used for bandwidth hungry applications like streaming high resolution videos and images. New applications continue to emerge in the form of wireless Internet of Things (IoT) devices such as home appliances, sensors, smart TVs and other media devices. Internet protocol (IP) based TV services such as Netflix, Hulu and Youtube have displaced conventional cable and satellite. Increase in display resolutions and video quality has caused streaming to consume a larger amount of bandwidth. Cellular technologies like 3G and 4G-LTE are subject to data-caps to deter over-consumption. Internet service providers offer high data-caps and unlimited packages are common and affordable. This has mitigated overloading of cellular networks and has been the primary motivator for mobile users to offload data usage to Wi-Fi networks when at home, public establishments, and the workplace. This shifts a larger burden of the growing wireless demand on Wi-Fi networks to supplant the scaling capacity requirements. However as with the tragedy of the commons, network over-utilization results in decreased overall performance for all users due to contention, congestion and interference.

1.2 Purpose

There is an ongoing race between wireless demand and capacity as the electromagnetic (EM) spectrum has become cluttered, and network performance is limited due to technological and scaling limitations. This has rendered old Wi-Fi technologies insufficient and created a demand for the next generation of wireless infrastructure enhancements. Very-High-Throughput (VHT) amendments have been proposed for WLAN in the form of 802.11ac (5 GHz) and 802.11ad (60 GHz). Technological difficulties associated with 802.11ac present significant design challenges, and also scaling limitations in practical deployments, such as:

- Interference: wider channel bandwidth reduces the total number of non-overlapping channels leading to increased congestion and interference with neighbouring networks
- Static Channel: higher order enhancements [e.g. multi-user multiple-input multiple-output (MU-MIMO)] require strong signal-to-noise ratios (SNR) which is only possible in very idealized channel conditions (e.g. stationary clients)
- Cost: 802.11n access points (APs) and clients must be upgraded to support 802.11ac which is significantly more expensive, and other devices in the backbone network (e.g controllers, edge switches, PoE adapters) must be upgraded[7]
- Chicken vs. Egg: It might take a long time for client side adoption to leverage infrastructure enhancements due to high transceiver complexity and cost

IEEE 802.11ad is another VHT amendment to the 802.11 family but in the unlicensed 60 GHz band. Accordingly, it faces a wide array of challenges and limitations unique to the 60 GHz band, such as:

- The oxygen absorption spectrum also coincides with the millimeter wave band (mm-wave) resulting in approximately 16db/km attenuation at sea level [8]; although this is more applicable for long range outdoor applications
- Increased free-space path loss (FSPL) at higher frequencies
- Lower reflection and higher material absorption
- Quasi-optical propagation: even worse signal penetration than 5 GHz signals but non-line-of-sight (NLOS) link can be maintained using reflections and lower order modulation techniques

1.3 Problem Statement

IEEE 802.11ad is a critical amendment to the WLAN family to augment upcoming generations of AP products and consumer devices to satisfy the increasing throughput demands over short ranges. Adding support for 60 GHz will alleviate congestion and contention on the longer range 2.4/5 GHz counterparts and increase overall WLAN capacity, throughput and reliability for all users. This thesis contains an in-depth performance and reliability characterization of early generation 802.11ad platforms. This included testing and evaluation of several commercial off-the-shelf (COTS) 802.11ad platforms and a custom 802.11ad proof-of-concept (PoC) implementation to characterize the efficacy of this new WLAN amendment and to identify the capabilities and limitations. Some of the important factors this research aims to address include:

- Quantify obtainable performance, stability and scalability in real world environments such as conference rooms, open office spaces, and cubical farms
- Establish spatial boundaries for supported single-carrier data rates in ideal line-of-sight (LOS) scenario [Rate vs. Range (RvR)]
- Qualify the effects of persistent and temporal obstructions such as doors and human blockages
- Determine the effects of station mobility and orientation on link quality
- Characterize scaling limitations associated with multiple-client (multiclient) spectral sharing by measuring throughput degradation as the number of contending nodes in a single network increases
- Feasibility analysis of frequency-reuse: determine the adverse effects of multiple overlapping networks [multi-AP (co-channel interference (CCI) and adjacent channel interference (ACI)], and determine how the effects change as network density increases
- Identify system engineering challenges associated with VHT support in the 60 GHz band such as host side processing limitations and radio frequency (RF) front-end design considerations such as antenna placement, orientation and topology
- Identify the overall strengths and weaknesses of the technology to dispel myths associated with 60 GHz Wi-Fi and to establish applications and use-cases

1.4 Thesis Organization

This thesis is organized as follows. Chapter 2 contains the necessary background of the IEEE 802.11 WLAN family, including the basic network architecture, medium access control (MAC) mechanisms, and a timeline of the key WLAN amendments. It then outlines the major features of 802.11ad in terms of the physical (PHY) and MAC sublayers, and added important technology enhancing features. Differences between traditional Wi-Fi and 802.11ad [also known as Wireless Gigabit (WiGig)] are explained and related to core components of the standard. Additionally Chapter 2 contains a summary of existing research pertinent to 802.11ad including link layer and system level simulations and theoretical research.

In Chapter 3 link budget analysis is conducted to create transmission models for idealized LOS wireless performance in the 60 GHz band based on SNR requirements for supported single-carrier data rates. These models and existing research are used to establish a performance and reliability baseline for the technology.

Chapter 4 outlines the experimental setup for real world 802.11ad testing including description of test platforms: APs, client stations (STAs) and wireless docks. Additionally, Chapter 4 outlines the PoC platform and the various antenna topologies and orientations that were tested. Lastly, Chapter 4 outlines the environments where 802.11ad platforms were tested and which real-world deployment scenarios they attempted to address.

In Chapter 5 the results of the test of each of the suites are presented along with detailed explanations and important observations.

Finally, in Chapter 6 the overall results are summarized and future work is proposed.

Chapter 2

Background and Related Work

In this section necessary background information pertinent to 802.11ad WLAN is summarized. A description of Wi-Fi architecture and fundamentals are provided in terms of the MAC and PHY sublayers. The timeline of the evolution of the 802.11 WLAN family is enumerated with emphasis on significant differences. After an introduction to the mm-wave frequency band, a detailed overview of 802.11ad MAC and PHY layers are provided and a summarization of important features unique to this WLAN amendment. Finally, existing research related to the 60 GHz WLAN amendment are summarized and categorized into either real world testing or studies on the efficacy of various MAC layer enhancements.

2.1 Wi-Fi Fundamentals

2.1.1 The 802.11 Network Architecture

The fundamental building block of the 802.11 network is the basic service set (BSS). A BSS contains a central base station, known as an AP, and one or more wireless STAs. The central AP would be connected to other networking devices and the Internet through devices such as switches and routers. Multiple wireless clients communicate with each-other by relaying messages through the central AP. This operating mode of 802.11 WLANs is known as infrastructure BSS; where infrastructure refers to the APs and the wired Ethernet backbone connecting the networks through routers. When multiple BSSs become connected through switches and routers they form an extended service set (ESS) connected through a distribution system. ESSs are denoted by a service set identifier (SSID). Home

routers typically combines both AP and routing functionality, and sometimes even act as 'gateways' by incorporating a DSL or cable modem. An alternative mode of WLAN is independent BSS (IBSS). IBSS is an ad hoc network with no central APs; meaning they cannot connect to other BSSs. In an ad hoc network, STAs communicate peer-to-peer (P2P) while obeying the MAC mechanisms. Ad hoc WLAN networks are less common but can be used to interconnect small cells of Wi-Fi devices without the need for a coordinating AP. In this thesis, focus is directed towards infrastructure wireless LANs although most of the theory applies to either operating mode.

2.1.2 Channels and Association

In 802.11 each station must associate with an AP before it can receive networked data or connect to the Internet. When stations are associated to an AP it is conceptually the same as attaching a virtual Ethernet cable between them. The 802.11 standard requires APs to periodically send beacon frames which contain its MAC address, SSID and other important information. In a passive scan, wireless clients scan the supported channels (e.g. 11/14 in the case of 2.4 GHz, 3/4 in the case of 60 GHz in North America), seeking beacon frames to learn about the available networks. Wireless stations can also conduct an active scan, where they broadcast probe request frames. 802.11 APs within range transmit probe response frames containing similar information to the beacon frame such as the SSID, MAC, supported data rates and encryption modes. Once the client has decided a network to join, it sends an association request frame to which the AP will respond with an association response frame either accepting or denying the client connection. After association clients must get either a static or dynamic IP address to begin communicating.

2.1.3 Medium Access Control Protocol

Access to the wireless medium in 802.11 networks requires a multiple access protocol in order to avoid collisions by coordinating transmissions. There are essentially three classes of multiple access protocols: channel partitioning (such as orthogonal frequency division multiplexing used by cellular standards), random access (such as with Ethernet), or turn-based (e.g. scheduled time-division multiplexing). The designers of 802.11 used a random access protocol due to its success in Ethernet protocol. With Ethernet, if a collision is detected mid-transmission, vying stations halt sending and wait different random backoff intervals before attempting to re-transmit. The idea is that the random backoff will establish a fair sequential order for retransmission. The random access protocol with WLAN is

known as carrier sense multiple access with collision avoidance, or CSMA/CA. The CSMA is the same as used in Ethernet. Carrier sensing is where each STA listens to the medium for transmissions [clear channel assessment (CCA)] and only attempts to transmit after a certain idle time. This is analogous to waiting until someone is finished talking before speaking up in a meeting. The difference with Ethernet is that it uses collision detection (CD) instead of the collision avoidance (CA) employed by WLAN. CA must be used for WLAN in lieu of CD due to some fundamental difficulties associated with wireless.

Ethernet is full-duplex, meaning an endpoint can both send and receive at the same time. When an Ethernet node transmits, it also listens for other transmissions. If it detects a collision the transmission is stopped immediately and the random backoff is initiated. A wireless medium is different as it exceedingly difficult to create a full-duplex wireless transceiver. The difficulty arises from self-interference. If a radio were both sending and listening on the same channel, its own transmission would drown out any background transmissions by having significantly higher signal power due to antenna co-location. A simple analogy is that in the wireless world, transmitters are shouting out as loud as they can, but at the distance of the receiver fading has caused the signal to diminish to a faint whisper. The receiver must then apply filtering and low-noise amplification to extract the original message. Furthering this analogy, it would be difficult to carry a conversation with someone across the room if another person were speaking loudly in your ear. For this reason an 802.11 device always transmits the frame in its entirety, hence if there is a failure in reception, it is much more costly than with Ethernet. In order to detect collisions, and to overcome reception failure due to the relatively higher bit-error rate on a wireless medium, WLAN standards include link layer acknowledgement (ACK) and retransmission control frames. If an ACK frame is not received for a transmission, the sender attempts to retransmit using the CSMA/CA algorithm. Full-duplex wireless active area of wireless research such as in [9, 10, 11].

Collision detection is infeasible in a wireless medium for a second reason known as the hidden terminal problem, pictured in Figure 2.1. Nodes A and C are client stations, and B is the AP. The circles indicate the range of each station due to fading. It is clear in this case that A,B are within each others range and so are B,C, however A,C are too distant to detect each others transmissions. If both A and C send packets to B simultaneously, they will collide at B resulting in failed reception. Alternatively, terminals may be hidden from each other due to physical blockages. In a hidden node scenario carrier sensing is ineffective since endpoints cannot sense each others transmission and will attempt to transmit resulting in a collision. To overcome this, 802.11 implements control frames request to send (RTS) and clear to send (CTS). Before transmitting, a STA sends an RTS and waits for a CTS granted by the central AP. This alleviates the hidden node problem as the AP will not

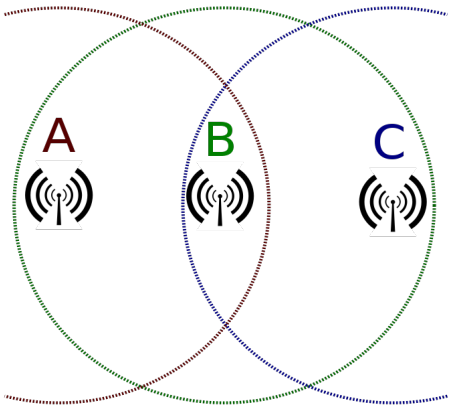

Figure 2.1: Hidden Terminal Diagram (Source [12])

issue CTS to multiple STAs with overlapping transmit opportunities (TXOPs).

2.1.4 The 802.11 Frame

The WLAN frame shares many similarities with an Ethernet frame, but has some additional fields that are specific to wireless links. A WLAN frames payload field can be up to 2312 bytes, but is typically limited to below a 1500 Ethernet maximum transmission unit. The cyclic redundancy check (CRC) field is used to detect bit errors which are more common in wireless WLAN than in wired.

The 802.11 frame also contains some important control fields. The sequence number is used to uniquely identify packets for easier acknowledgment of reception and retransmission. Lastly, fields type and subtype play an important role for distinguishing management frames such as association, control frames such as RTS, CTS, ACK, and data frames. The duration field, is used to indicate the channel reservation time for a TXOP.

A significant difference in the WLAN frame is the inclusion of four address fields as opposed to the two address fields in a conventional Ethernet packet. The first address is the destination MAC address of the wireless station. The second is the source MAC address of the wireless station. The extra third address is used to indicate the MAC address of a layer three routing interface for crossing between subnets. Essentially this additional address is for connecting a BSS with a wired LAN [13]. This address is important because the router will have to convert the 802.11 frame to the 802.3 Ethernet frame, with the destination MAC address being the third address of the original 802.11 frame. The fourth address is only used for wireless bridges. There are many scenarios where address fields remain

unused. For example, with CTS frames, only the source address is mandatory. Another example is that control and management frames do not include the fourth address whereas a data frame could if the destination is across a wireless bridge.

2.1.5 Advanced Features

Rate Adaptation

An important feature of the WLAN family is rate adaptation. WLAN standards support multiple data-rates with varying SNR requirements. This allows robust communication in a wide variety of scenarios. Network interface cards implement their own proprietary algorithms for rate-shifting. Typically an initial rate is selected based on the received signal power of incoming packets. After a certain threshold of reception failures the rate may downgrade to a lower throughput but more reliable modulation scheme. Conversely, in the absence of interference after a certain threshold of successful transmissions the sender may upgrade to a higher throughput modulation scheme. Rate adaptation is an important feature in WLAN which allows stations to communicate at varying ranges and in the presence of dynamic channel conditions.

Power Management

Another important consideration for electronics is power consumption. Energy efficiency is of paramount importance to extend battery life of mobile devices such as laptops and smartphones. 802.11 devices support enhancements which facilitate power saving modes such as sleep and doze states which minimize the time their sense, transmit and receive circuitry must be active. Stations can announce when they enter power saving modes by setting a flag in the 802.11 header and the AP will buffer transmissions for that node until the station wakes. The station can then set a sleep timer to wakeup sometime in the future before the AP transmits its beacon. During its beacon the AP indicates which nodes have buffered transmissions so it can wake, check the beacon to know if it missed any packets, then it can re-enter sleep state. Additionally, various data rates require ranging complexity of radio hardware with diverse power footprints. Only subsets of data rates are mandatory so portable devices can support energy efficient rates that satisfy their application demands. These features enable fine tuning of power footprints which contributes to the success of the WLAN family.

2.2 The Evolution of Wireless LANs (IEEE 802.11)

Due to changing demand and application needs the IEEE 802.11 WLAN standard has been amended every few years to add features which achieve different target data rates and ranges through use of various modulation and encoding techniques and frequency bands. Iterative improvements in the WLAN family require upgrading of wireless devices to leverage new technological advancements. The sections herein outline WLAN technologies and summarize the most relevant features of each amendment.

802.11-Legacy (1997)

The original draft of the WLAN family operated in the 2.4 GHz Industrial Scientific Medical (ISM) frequency band. The original standard defined three PHY layer technologies: diffused infrared, frequency-hopping spread spectrum (FHSS) and direct sequence spread spectrum (DSSS) and enabled data rates of 1 and 2 Mbps.

802.11b (1999)

Due to high resistance to narrowband interference 802.11 with DSSS eventually replaced other PHY technologies and was popularized in 802.11b. 802.11b increases the maximum data rate to 11 Mbps by supplementing Barker codes with complementary code keying for modulation but ultimately suffered from increased interference with other devices operating in the 2.4 GHz band such as Bluetooth, cordless phones and microwave ovens. The unprecedented throughput and price reduction offered by 802.11b ultimately lead to widespread WLAN adoption.

802.11a (1999)

802.11a was the first LAN technology which uses orthogonal frequency division multiplexing (OFDM) for modulation instead of DSSS and operated in the 5.8 GHz band. The use of the relatively unused 5.8 GHz band allowed for less interference and better throughput, but due to a higher carrier frequency it suffered from higher signal attenuation and absorption loss than its 2.4 GHz counterparts. The standard supports up to 54 Mbps throughput due to the use of 52 data sub-carriers and higher complexity OFDM.

802.11g (2003)

802.11g matched the 54 Mbps offered by 802.11a in the 5 GHz band and brought it to the 2.4 GHz band through the use of OFDM instead of DSSS. At the cost of higher transceiver complexity, OFDM enabled higher throughput in both bands due to more bits being encoded per symbol, while still spreading the signal over the frequency range and allowing resilience against frequency-selective fading and narrowband interference. 802.11g remained backwards compatible with 802.11b which caused downgrade of overall network performance by approximately 20% in the presence of any 802.11b clients [14]. Similar to 802.11b, 802.11g suffered from interference with other devices in the 2.4 GHz band.

802.11n (2009)

The 802.11n WLAN amendment greatly improved achievable throughput by adding support for single-user multiple-input multiple-output (SU-MIMO). 802.11n operates in both 2.4/5 GHz but 5 GHz support is optional. SU-MIMO takes a single data stream and divides it up into multiple data streams which can be sent from multiple transmit antennas to multiple receive antennas (on the same client) simultaneously. The spatial streams are re-combined at the receiver increasing the overall throughput at the cost of higher transceiver complexity and more antennas. The maximum number of spatial streams is four but is ultimately limited to the number of client antennas. 802.11n added maximum data rates from 54 Mbps to 600 Mbps.

802.11ac (2013)

802.11ac is a 5 GHz-only technology which improves on the older 802.11n amendment by adding wider channels (80, 80 + 80 bonded, and 160 MHz), multiple-user MIMO (MU-MIMO), an enhancement to MIMO where multiple spatial streams could be sent to multiple clients concurrently. This would mitigate much of the degradation effects associated with high contention for channel access between multiple 802.11ac clients. 802.11ac added support for up to 8 spatial streams, and higher order modulation [256 Quadrature Amplitude Modulation (QAM) vs. 64 QAM of 802.11n]. Due to the complexity of features added by 802.11ac deployment was separated into two technological iterations: Wave I and Wave II. Wave I devices offered a maximum data-rate 433 Mbps per spatial stream with high end implementations supporting up to three spatial streams; a total throughput of 1.3 Gbps. Even then, client devices could rarely see such high performance due to the high SNR

requirements. Wave II added 160 MHz channels and MU-MIMO support however technological challenges have caused limited support for all these features even in enterprise-grade products.

2.3 Millimetre Wave Communication

The mm-wave band can be used for high speed wireless broadband communication and lies within the extremely high frequency (EHF) and very high frequency (VHF) bands denoted by the International Telecommunications Union (ITU). Mm-wave band wavelengths range from 1-10 mm with frequencies between from 30-300 GHz. Signals in this frequency range suffer from increased attenuation by being absorbed by gases in the atmosphere as well as reduction in signal strength due to rain and humidity [15]. Several applications of mm-wave technologies have emerged in recent years as a result of the lack of bandwidth or availability for some of the lower frequency bands. High-throughput point-to-point links are used in the mm-wave ranging from 71-76 GHz, 81-86 GHz, and 92-95 GHz but require a license from the Federal Communications Commission (FCC). IEEE 802.11ad defines VHT in the unlicensed 60 GHz band for the popular WLAN family. IEEE 802.15.3c defines a wireless personal area network (WPAN) mm-wave PHY supporting up to 5.28 Gbps, and ECMA International also published a 60 GHz PHY, MAC and HDMI protocol adaptation layer (PAL) also for WPAN applications [14]. The emergence of these mm-wave technologies indicates a trend towards augmenting existing broadband wireless with small-cell VHT in the mm-wave band.

2.4 IEEE 802.11ad (2012)

Wireless communication in the 60 GHz band has different characteristics compared to its 2.4/5 GHz counterparts. The following sections contain background information on the 802.11ad amendment in terms of the physical layer, MAC mechanisms and the key technology enhancing features.

2.4.1 Physical Layer (PHY)

Channelization

In 2001, the US FCC allocated 7 GHz in the 57-64 GHz frequency range for license-free usage[16]. The global spectral allocation for 802.11ad is from 57-66 GHz. The spectrum allocation is divided into four 2.16 GHz non-overlapping channels with center frequencies shown in Figure 2.3. Channel support varies by region but most support at least two channels. IEEE 802.11aj is a re-banding of 802.11ad to the 45 GHz band primarily for unlicensed use in China [17]. Since channel 2 (F_c=60.48 GHz) is supported in all regions it is the default channel used by early technology vendors. Figure 2.2 shows the transmit mask defined for 802.11ad in order to mitigate adjacent channel interference caused by overlapping transmissions.

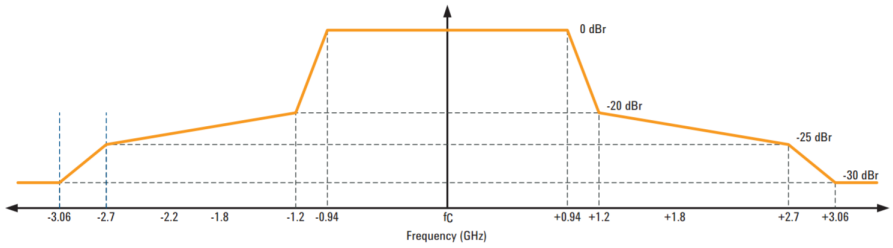

Figure 2.2: 802.11ad Transmit Mask (Source [18])

A channels maximum capacity is governed by the Shannon-Hartley equation shown below:

$$C = BW \log_2(1 + SNR) \tag{2.1}$$

where:

C = The channel capacity (maximum transmission throughput) in bits per second
BW = The channel bandwidth (Hz)
SNR = The signal to noise power ratio (W/W)

One important implication of the Shannon-Hartley relationship is that a channels maximum bit-rate is directly proportional to the channel bandwidth. A lower sampling frequency limits the fastest transition that can be captured, therefore the bandwidth limits the maximum bit-rate of the data signal that can be modulated into the carrier. Since the bandwidth of an 802.11ad channel is 2.16 GHz, compared to the 20 MHz channel

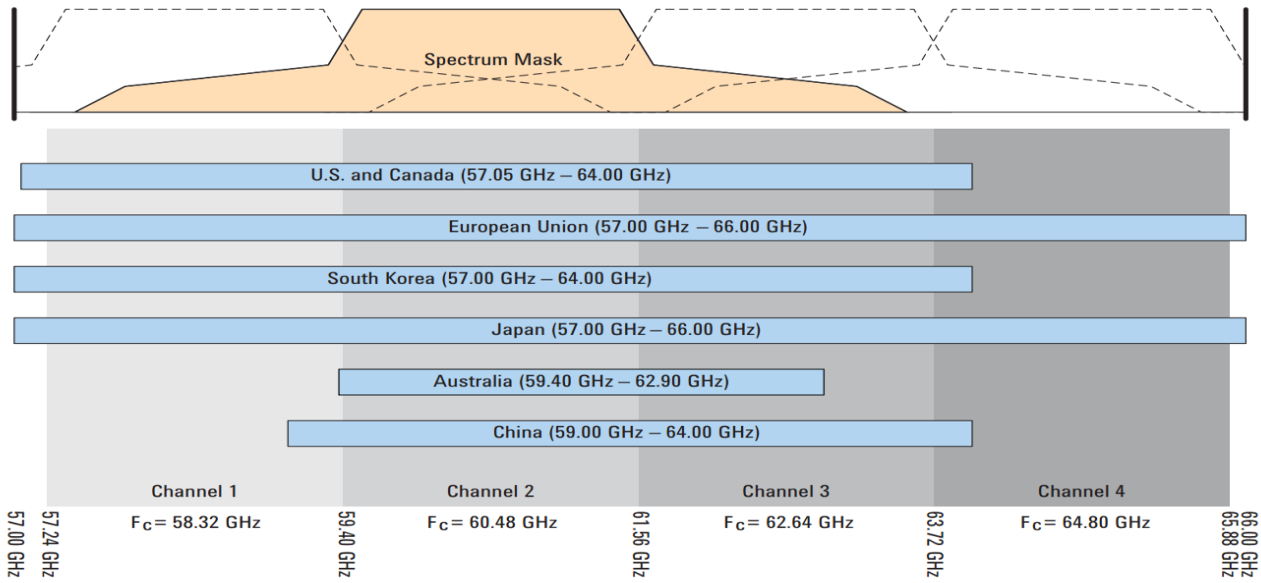

Figure 2.3: 802.11ad Channelization By Region (Source [18])

bandwidth of a traditional Wi-Fi channel, the achievable capacity is inherently higher by a factor of one hundred with all other factors being equal. In practice the same SNR cannot be obtained in the 60 GHz channel for a variety of reasons. Due to the channel bandwidth being wider, there is also significantly more noise that enters the RF hardware increasing the denominator of the SNR. Additionally, the larger FSPL and material absorption at 60 GHz results in a lower signal power level due to fading, decreasing the numerator of the SNR.

IEEE 802.11n was the first WLAN standard to introduce channel bonding allowing two 20 MHz channels to be aggregated for an effective channel width of 40 MHz. This essentially doubles the throughput, but can suffer from increased interference since two non-overlapping channels are required. The 5 GHz VHT amendment 802.11ac supports even wider 80 MHz and 160 MHz channels. Due to the high bandwidth of WiGig and associated complexity of development, advanced wireless enhancements such as channel bonding were excluded from the standard. IEEE 802.11ay is the next generation of 802.11ad and proposes data rates up to 20 Gbps for point-to-point by leveraging enhancements such as MU-MIMO, channel bonding, and increased range up to several hundreds of meters [19]. The IEEE has identified that the 60 GHz spectrum is very promising for VHT applications such as cellular offload, backhaul, WPAN, and augmented/virtual reality (AR/VR). Since 60 GHz WiGig is something that will be invested in heavily and enhanced over time, research into

all aspects of this technology is important to demystify practical performance and identify applicable use cases.

Modulation Coding Schemes

IEEE 802.11ad standard defines three heterogeneous PHYs for different operating modes. All packets share a common preamble modulated with BPSK to ensure proper link training and synchronization followed by a header and payload with different MCS depending on the PHY in use. The first is the control PHY (CPHY). CPHY uses simple DBSPK to increase range and reliability at the cost of lower throughput. Control and management frames such as beacons and probes are sent using the CPHY to ensure successful reception. The second is the single-carrier PHY (SCPHY). SCPHY uses more advanced modulation schemes up to 16-QAM supporting throughput up to 4.6 Gbps. There is an additional variant of the SCPHY which uses Reed-Solomon (RS) codes for encoding as opposed to a low-density parity check (LDPC). Use of RS encoding enables lower power consumption which makes it an ideal coding scheme for battery powered mobile devices. Other than the encoding, the modulation-coding scheme (MCS) is identical to SCPHY so it is simply renamed low-powered single-carrier PHY (LPSCPHY). The third PHY variant is the orthogonal frequency division multiplexing PHY (OFDMPHY). The OFDMPHY provides data rates up to 6.75 Gbps over 1.825 GHz of bandwidth. OFDMPHY uses 512 subcarriers total (336 data and 157 null) and uses up to 64-QAM to provide VHT with resistance to frequency selective fading. There are a few MCS rates for each PHY which have mandatory support which provides a trade-off between universal interoperability, and finely tuned power-performance profiles. The supported coding, modulation types and bit rates for each PHY is summarized in Figure 2.4.

Packet Structure

Each PHY packet structure is shown in Figure 2.5. After the common preamble modulated with BPSK the header and payload modulation varies depending on the employed scheme. The first field is a short training field (STF) is used for timing estimation, and to adjust automatic gain control circuitry on the radio front end. Essentially it is used for synchronization and to ensure voltage levels are of correct magnitude for the low noise amplification circuitry. The second field is a channel estimation field (CEF). The next field is the physical layer convergence protocol (PLCP) header which contains some fields such as the packet length, training length [optional beamform training (BFT) at the end], and

Control (CPHY)		
Coding	Modulation	Raw Bit Rate
Shortened 3/4 LDPC, 32x Spreading	π/2-DBPSK	27.5 Mbps
Single Carrier (SCPHY)		
Coding	Modulation	Raw Bit Rate
1/2 LDPC, 2x repetition 1/2 LDPC, 5/8 LDPC 3/4 LDPC 13/16 LDPC	π/2-BPSK, π/2-QPSK, π/2-16QAM	385 Mbps to 4620 Mbps
Orthogonal Frequency Division Multiplex (OFDMPHY)		
Coding	Modulation	Raw Bit Rate
1/2 LDPC, 5/8 LDPC 3/4 LDPC 13/16 LDPC	OFDM-SQPSK OFDM-QPSK (DCM) OFDM-16QAM OFDM-64QAM	693 Mbps to 6756.75 Mbps
Low-Power Single Carrier (LPSCPHY)		
Coding	Modulation	Raw Bit Rate
RS(224,208) + Block Code(16/12/9/8,8)	π/2-BPSK, π/2-QPSK	625.6 Mbps to 2503 Mbps

Figure 2.4: 802.11ad Modulation Coding Schemes (Source [18])

lastly received signal strength indicator (RSSI) and MCS. The PLCP also contains an intial value for the receive scrambler (SCPHY), or the subcarrier initialization (OFDM). The next field is the payload (MAC frame), a CRC, followed by some additional BFT fields such as another automatic gain control field and other training fields. IEEE 802.11ad makes extensive use of heavy spreading and Golay complementary sequences G_{32}, G_{64}, Ga_{128} and Gb_{128} in the STF, CEF and PLCP due to their useful auto-correlation properties which allow robust signal decoding even in the presence of noise and interference [8].

2.4.2 Medium Access Control Layer (MAC)

Hybrid MAC Layer

802.11ad utilizes a hybrid MAC protocol to coordinate directional communication between several nodes. Through the use of a combination of deterministic time-division multiple access (TDMA) service periods (SPs) and traditional CSMA/CA access periods 802.11ad facilitates scheduled SPs, dynamic SPs and CBAPs. In P2P mode (IBSS) stations provide periodic interference reports regarding any other transmissions they detect. An AP can

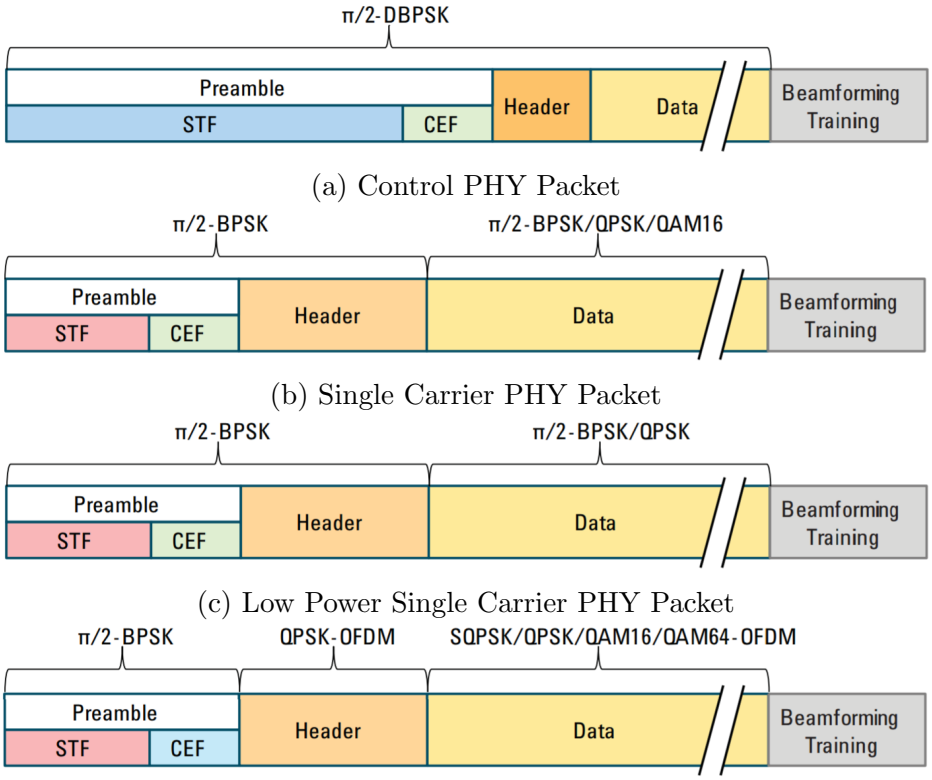

(a) Control PHY Packet

(b) Single Carrier PHY Packet

(c) Low Power Single Carrier PHY Packet

(d) OFDM PHY Packet

Figure 2.5: IEEE 802.11ad PHY Packet Structures (Source [18])

use this information to grant overlapping TXOPs between non-interfering stations. This is known as spatial sharing, a feature entirely unique to 802.11ad due to the high frequency re-use offered by the mm-wave band. 802.11ad supports SP allocation, truncation and extension to support dynamic load profiles.

When an 802.11ad node wants to transmit, it obeys the hybrid TDMA-CSMA/CA protocol based on 802.11 enhanced distributed channel access (EDCA). EDCA provides two carrier sensing mechanisms. The first is the physical CCA whereby the node senses the medium to detect transmissions. The second is a virtual sensing mechanism. When an 802.11 node detects a CTS intended for another station it sets a timer known as a network allocation vector (NAV) to indicate the time (in microseconds) that the medium will be busy for. The NAV duration is indicated in a header field of the CTS frame. When both the physical and virtual sensing mechanisms detect the medium as idle, an 802.11 station can request to transmit by sending an RTS to the AP.

Beacon Interval

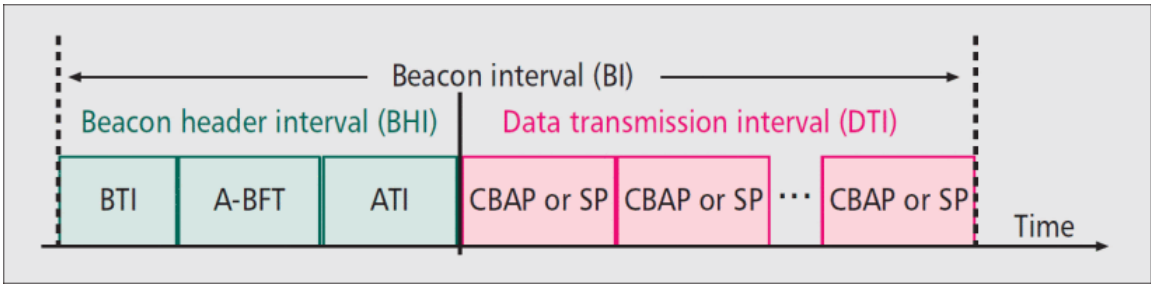

Figure 2.6: 802.11ad Beacon Interval (Source [20])

Channel access time in the 802.11 family is completed during the beacon interval (BI). In 802.11ad the BI comprises of four sections as shown in Figure 2.6, namely the beacon transmission interval (BTI), association beamform training (A-BFT), announcement transmission interval (ATI), and data transmission interval (DTI). In the BTI, the AP conducts a sector sweep (SSW) and broadcasts beacon frames in all spatial directions. In the following A-BFT slot, associated stations provide explicit feedback (SSW-FB) about which training packets from the BTI were received with the highest signal strength. This allows the AP to maintain the spatial direction of each associated station. In the ATI, the AP announces the schedule for data transmission in the subsequent DTI. A mix of scheduled SPs and CBAPs are allocated in orthogonal time slots (or overlapping time slots permitted spatial sharing is coordinated between non-interfering stations).

Frame Structure

The 802.11ad MAC frame is pictured in Figure 2.7. The MAC frame consists of both overhead fields (header) and data (payload). In 802.11 parlance the payload is referred to as the service data unit (SDU). The overhead combined with the SDU together are referred to as a protocol data unit (PDU). Each MAC frame is suffixed with a frame check sequence (FCS) which is used as an error-detecting code. When an MSDU is received the manually computed FCS is compared with the attached FCS. If there is a discrepancy the frame is discarded. The encrypted MPDU, shown in green, notably includes an 8-bit Galois Counter Mode (GCM) protocol header and a 16-byte message integrity check to detect malicious modifications that may slip past the FCS. Compared to the size of a typical payload (thousands of bytes), the overhead due to these fields does not significantly degrade throughput.

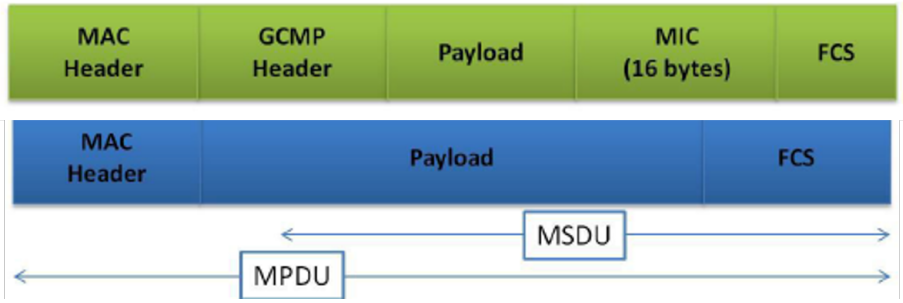

Figure 2.7: 802.11ad MAC Frame Structure (Source [21])

2.4.3 Important Features

Security

Security in 802.11ad is similar to previous WLAN standards. It still supports Advanced Encryption Standard (AES) but deploys a different operating mode known as GCM. AES is a block-cipher thus a plaintext message is divided into 16-byte blocks. Traditional Wi-Fi uses the Cipher-Block-Chaining (CBC) Counter Mode (CCM) of AES. With CCM when each block is encrypted the enciphered output is chained into the next block encryption, hence the block-chaining nomenclature. This ultimately enhances security but is unsuited for multi-Gbps communication due to the data dependencies between subsequent ciphertext blocks. An Ethernet maximum transmission unit (MTU) is 1500 bytes and each block requires two rounds of AES. Therefore encryption of an MTU requires roughly 200 serial AES operations with CCM. This serialization is a bottleneck for the VHT offered by 802.11ad. GCM requires only one AES operation per block, and removes the block chaining allowing parallel computation of block ciphertexts. The removal of data dependencies enables multiple ciphertext blocks to be computed simultaneously; allowing for hardware and software parallelization. GCM has been proven secure in the concrete security model [22]. With any mode of AES, secure operation requires proper usage of unique initialization vectors for every encryption with the same key to prevent a stream cipher attack.

There is another interesting aspect of 802.11ad security at the physical layer. Due to the reduced signal penetration at 60 GHz and reliance on highly directional beamformed links the feasibility of eavesdropping is greatly reduced. With traditional Wi-Fi bands it is possible to intercept 2.4 GHz 802.11n packets from as far as 820 ft away[23]. With 60 GHz this is extremely impractical which is why the mm-wave band was initially suggested for battlefield communications[24]. This also allows for high frequency re-use, where neighboring networks can operate on coincident or adjacent channels without significantly degrading

each others performance and reliability. This is a limitation with traditional Wi-Fi deployments as other networks on the same channel must take turns accessing the medium, and adjacent channel communications are received as increased noise resulting in lower SNR. This has sparked advanced studies into radio resource management (RRM). RRM involves strategies and algorithms for increasing spectral efficiency by controlling transmit power to reduce service area overlap and mitigate adverse effects of multi-cell networks. RRM is an ongoing research area and comprises of both static and dynamic techniques. Transmit power control is an example of a dynamic RRM technique. In summary, WiGig's weakness is also a strength in terms of frequency re-use, system spectral efficiency, and physical layer security.

MIMO Beamforming

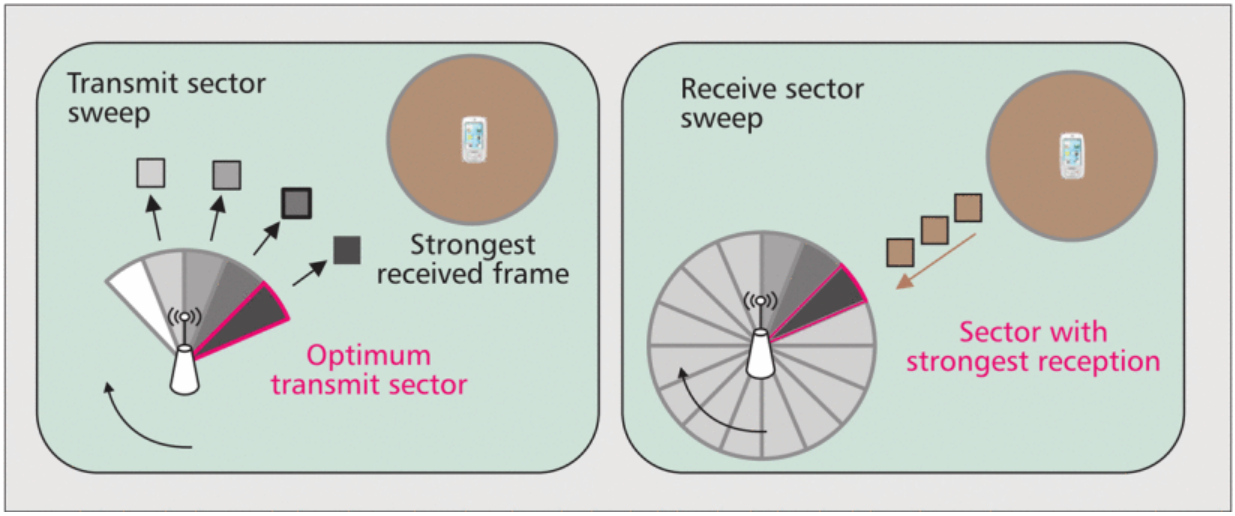

Figure 2.8: Beamforming Diagram (Source [20])

One of the primary enhancements required for DMG communication in the mm-wave band is beamforming. Due to the high density of mm-wave antenna elements WiGig front ends are designed using phased antenna arrays (multiple antennas) to transmit the same signal with slightly different phase offset such that the signals constructively interfere at the receiver. This facilitates increased SNR despite the higher path loss and attenuation experienced at 60 GHz. Beamforming in WiGig consists of three phases: i) sector sweep, ii) beam refinement, and iii) beam tracking. In order to compute the necessary phase and amplitudes (antenna weighting vector) BFT is conducted. BFT is within the BTI where

the AP will conduct a SSW by transmitting its beacon from each sector. The receiver will then respond with a SSW-FB indicating which packets were received with the highest SNR. Secondly, there is an A-BFT wherein stations will train their antennas for uplink communication to the AP. Beamforming is optional in the standard but all endpoints must support the SSW-FB protocol so that the AP can beamform on the downlink. Due to the high density of mm-wave antenna elements, it is highly feasible to see large phased antenna arrays on the client side however they might activate less concurrent transmit chains as a power-performance trade-off. A simple real-world analogy to beamforming can be made by imagining two people with flashlights in a large dark field. SSW is analogous to one person spinning in a circle with their flashlight on. When the flashlight is correctly aligned the second person shouts out indicating that the beam is in the correct spatial direction (SSW-FB). Figure 2.8 shows an example of the transmit SSW, and also how multiple receive antennas provide selection diversity.

Access Categorization

Similar to previous WLAN technologies, 802.11ad has several mechanisms to ensure load balancing and quality of service (QoS) between competing applications and client devices. Access to the 60 GHz channel is divided into four access categories (ACs) with varying user priorities from lowest to highest: background, best-effort, video and voice. This prioritization is accomplished by hosts having separate queues for data streams of differing priority. Then when a scheduled window or a CBAP has been granted, the host system will empty the queues following the aforementioned affinities. This ensures that latency critical data streams such as uncompressed video and voice-over IP are less likely to encounter a starvation scenario. In addition to having separate buffer prioritizations, the standard adjusts the contention window (CW) bounds based on the application demands. When multiple STAs vie for channel access in a shared wireless medium and a collision occurs, both STAs must set a random backoff timer wherein they will not try to access the medium again. The random value is generated between a CW minimum and maximum (i.e. $T_{backoff} = rand[CW_{min}, CW_{max}]$). This allows for another type of prioritization whereby different ACs are given different values for CW_{min} and CW_{max}. When a collision occurs a critical voice AC could be given a deterministic backoff of 1 time-unit. This increases channel efficiency by reducing the silent period following any collision. In the case of repeated collisions, the contention window minimum and maximum are doubled in an exponential backoff reducing the likelihood of another collision.

Power-Save Protocol

The standard also has a few features which increase feasibility for battery powered devices to support WiGig. Similar to existing 802.11 standards wireless stations can enter sleep and doze states when they do not need channel access which can increase battery life and make tri-band support more practical. The central AP is notified when STAs enter sleep states and can buffer packets for that station until a scheduled wakeup. The combination of power saving states, along with the low-power variant of the SCPHY allows increased power savings during channel usage, and during periods of inactivity. Features such as this allow WiGig to have even broader applications ranging from high throughput streaming with OFDMPHY, to portable devices using LPSCPHY and power-save protocol.

Fast Session Transfer

Fast session transfer (FST) is a mechanism where STAs and APs can negotiate supported frequency bands upon association, and either endpoint can initiate transfer to a different band. This could allow the link to be transferred to a slower, more reliable band (e.g. 2.4 GHz) in the presence of signal degradation or high channel contention. Conversely, FST can allow upgrading of link speed by detecting when it would be more optimal to use a VHT channel in 5 GHz or 60 GHz bands. This could be useful when a mobile STA has entered the coverage area of a WiGig enabled AP. The FST mechanism allows joint management of multiple bands, improved load balancing and optimization of wireless performance. FST is broken into two modes, transparent (where the MAC addresses of each radio are the same), and non-transparent (where each radio has a unique MAC address). While non-transparent FST has lower level implementation repercussions to security, FST is intended to be seamless to the end user and simply enhances performance and reliability. WiGig can also inadvertently improve performance for 2.4/5 GHz only clients. By virtue of offloading the most bandwidth hungry small-cell applications (e.g. video streaming) to the 60 GHz band channel access opportunities for conventional Wi-Fi clients are more abundant, and contention is reduced on the 2.4/5 GHz bands.

Protocol Adaptation Layer

Another important feature in the 802.11ad standard is the protocol adaptation layer (PAL). Due to the VHT offered by WiGig the potential applications skyrocket. The bandwidth is enough to enable wireless bus extension (WBE) such as USB3.0 and PCIe over wireless. The PAL is a hardware abstraction layer that converts the low level encapsulation of

802.11ad traffic into the format expected by a USB or PCIe device. This allows existing devices to operate over wireless with lowered re-engineering costs associated with driver development and system integration.

Packet Aggregation

802.11ad defines aggregation at both the PHY and MAC layers. With MAC packet aggregation, multiple MSDUs can be aggregated after one header; when the source and destination MAC addresses are common. MPDU aggregation is also supported. This results in less overhead for successive transmissions between the same endpoints as there is only one PHY overhead to multiple MPDUs. 802.11ad also introduced aggregated PHY PDUs (PPDU) and SDUs (PSDU). This allows a sender to transmit PPDUs/PSDUs back-to-back with one header and with no interframe spacing (IFS) or individual ACKs. The receiver can then send a block ACK with a bitmap indicating which PPDUs/PSDUs were successfully received. PHY packet aggregation is a unique feature to 802.11ad that can facilitate higher throughput burst transmissions and allow endpoints to fully leverage their channel access intervals. The older 802.11n standard introduced the reduced IFS (RIFS) of 2 microseconds compared to the 10 microseconds in 2.4 GHz band or 16 microseconds in the 5 GHz band. Even with the savings from using the RIFS transmitting two frames still required two full headers and two PLCP frames. Due to this it is still much more efficient to just send aggregated frames, which is why 802.11ac and 802.11ad do not use RIFS[25].

2.5 Literature Review

Due to the standard being ratified in 2012 there has been various research efforts on the performance and reliability of 60 GHz wireless networks. To date, most studies have been focused on computer simulation and simple test cases since commercial platforms haven't become available until 2016. Even with the first consumer tri-band routers being released by TP-Link, Netgear and Acelink earlier in 2016 [26, 27, 28], the FCC has only just begun certifying the first client devices [29, 30]. Interoperability is still a lacking as vendors refine specific features and standards compliance. Intel also has shown significant interest in the 60 GHz front but primarily for WPAN applications[31]. Up until now, practical 802.11ad testing mainly has been conducted by technology manufacturers such as Intel and Qualcomm. Most of the real-world testing is done in-house to develop features and understand the important design principles. The sections below categorize available 60 GHz research into two main categories:

- Real World Testing: Involving test platforms/evaluation kits to measure the link-level performance and feasibility of 802.11ad for various practical application domains
- Simulation and Theoretical Analyses: Use mathematical simulations and modeling of the 60 GHz PHY and MAC layers to identify the benefit of a proposed theoretical enhancement to the 60 GHz MAC layer

2.5.1 Real World Testing

Practical 60 GHz testing is most relevant to this thesis as almost all of the observations and discussion are based off of empirical results from realistic network configurations. The theoretical analysis chapter applies RF theory and conducts link-budget analysis to estimate 60 GHz point-to-point LOS performance which is directly comparable to the long hall RvR testing. However, realistic deployments are too complex in nature to model with simulation or mathematical analysis. The existing studies involving real-world testing of 60 GHz platforms are limited in that they were targeted to an application domain. Results from specific use cases are not generally applicable to most networking configurations because of their simplicity in nature such as testing only a single LOS point-to-point link or using an early generation systems with lacking support for many of the standard features. The research in [32] studies 802.11ad performance in inter and intra vehicle single hop deployment scenarios. The aim is to try to identify the efficacy at 802.11ad providing in-car wireless streaming for entertainment and multimedia. The evaluation kit they tested is not intended for general networking configurations and the real-world testing they conducted is tailored to vehicular applications. The research in [33] uses a 60 GHz software defined radio platform with a mechanically steerable horn antenna to profile a 60 GHz link at the physical layer. This research provides valuable insight into the 60 GHz channel but its applicability to 802.11ad is limited since there is no beamforming and the bandwidth of the prototype is a major bottleneck to the higher MCS rates. The research in [34] uses 60 GHz links on top-of-rack switches in a data center. Using dynamic traffic monitoring they use the 60 GHz links as wireless flyways. The study proves 60 GHz applicability to improve load balancing in this unique application domain but again is not applicable to general network topologies. In this thesis several 60 GHz platforms were tested in realistic office and home deployment configurations. The emphasis is on the capabilities and limitations of 802.11ad in real-world scenarios such as: long hall LOS RvR, environmental testing (open office, meeting room, NLOS cubical farm), multiclient testing, spatial re-use (multi-AP), obstruction and mobility testing. These test results provide insight into 60 GHz propagation behavior, degradation due to increased load and network

density, and finally the effects of station distance, orientation, mobility and blockage. The performance and reliability measurements in these real world tests can be used to predict the efficacy of 60 GHz technology in many application domains and environments. Lastly, a large component of this research involved engineering the custom PoC EM to identify the system design engineering challenges due to the massive processing bandwidth required by this leading edge WLAN technology and to study the measurable effects of antenna topology, placement and orientation on the provided service area.

Vehicle-To-X Feasibility Analysis

Researchers in [32] studied the PHY layer of 802.11ad and its feasibility for Vehicle-to-Device and Vehicle-to-Infrastructure (V2X) communication. They studied the link range, data throughput, packet error rate, and coverage capability. The link range and performance were studied using mathematical models and MATLAB simulations. The practical communication performance at the PHY and MAC layers were measured in vehicle using 802.11ad evaluation kits. Their results show high feasibility for 802.11ad to be used in V2X applications, and the test outcomes outline some of the practical challenges faced in 60 GHz due to large path loss, effects of human body/obstructions, and weak reflection gain. This research contains a very effective overview of 802.11ad for this application and identify that more research is needed to demystify the theoretical limits of this technology and the performance and reliability that is achievable in practice. Their test platforms were less refined than the reference designs used in this thesis, but also provided some helpful link layer statistics such as packet error rate and signal strength. They also investigate the importance of RF design considerations such as antenna placements for in-vehicle applications.

Wi-Mi Software Defined Radio

Researchers in [33] released testing in 2015 which contains a detailed link level profiling of a customized 60 GHz software defined radio platform. They identified that to date, evaluation of 60 GHz protocols relies mostly on simulation complemented by analytical propagation models. Due to the complex nature of diverse beam patterns and sophisticated multipath reflections, it is virtually infeasible to reproduce mm-wave channel profiles with a unified model[33]. In the presence of obstructions and mobility, simulation and theoretical analysis are even more primitive at estimating performance in real world environments. The first research objective was to qualify the coverage and bit-rate of a single link. Their second objective was to identify the impact of beam steering on network performance

under human blockage and device mobility. Lastly, they investigated the feasibility of spatial re-use between flexible beams.

Augmenting Data Center Networks with 60 GHz

Researchers in [34] investigated usage model feasibility of DMG 60 GHz links in oversubscribed data center (DC) networks. They accomplished this by adding 60 GHz links on top-of-rack (ToR) switches to form wireless flyways for the wired DC network. They presented a design that uses DC traffic levels to select and add flyways to the wired network. By evaluating network traces they showed that predictable traffic workloads on a 1:2 oversubscribed network can be sped up by 45% in 95% of the cases with one wireless device per ToR switch. With two devices in 40% of the cases the performance was no different from an non-oversubscribed network. Their research investigates the usage of various antenna styles, the effects of human obstructions, and the feasibility of running concurrent 60 GHz links. This is one of many interesting and unique applications of 60 GHz and this research reinforces that the massive throughput offered by new 60 GHz technologies can offer great improvements to not just home networks but many real-world application domains.

2.5.2 Simulation, Modeling and Theoretical Enhancements

802.11ad Simulation Model in ns-3

Researchers in [35] presented a simulation model for 802.11ad using the network simulator ns-3. They model new techniques essential for 60 GHz operation such as beam steering, relay support, and FST. They showed that a pair of DMG endpoints with 8 sectors spent approximately 572 us to complete transmit SSW/SSW-FB and 396 us to complete receive SSW/SSW-FB during beamforming refinement. This result shows that beam alignment overhead was less than one millisecond which is suitable for latency critical applications (e.g. VoIP, VR/AR), and also allows DMG in the presence of station mobility (at low speeds). They also showed that the highest achievable TP for SC rates was almost 4 Gbps, and OFDM rates up to 5.2 Gbps which is 1.5 Gbps less than the theoretical maximum of the standard due to overhead imposed by the CBAP mechanism. In their relay simulations they demonstrate that in the presence of blockages OFDM TP quickly degraded from 5 Gbps to less than 2.5 Gbps without relay, whereas a relay station could regain 5 Gbps TP within 117 us. Lastly, their simulations showd that once a blockage was inserted FST mechanism allowed fallback to the 2.4 Ghz band resulted in a significant TP degradation (5 Gbps to 60 Mbps) but prevented loss or interruption of service.

Intelligent Listening During Association Beamforming Training (ILA)

Researchers in [36] proposed an enhancement to the original 802.11ad MAC design which would reduce beamforming training overhead during STA to STA communication in large networks with a numerous (>10) concurrent flows. They designed an intelligent listening mechanism where STAs listen to beamforming training between other STAs and the AP during the A-BFT of the BHI. By listening to other STAs sector sweep with the AP, they can get their approximate beam direction. This can considerably reduce the beamforming overhead during STA-STA communication by reducing the number of transmissions in BFT which results in less collisions for other transmitting nodes. They developed a set of MATLAB simulations where between ten and twenty STAs are placed within a 100 m^2 room. Each STA randomly selects another STA it wants to communicate with. What they demonstrated is that as the network scaled and more nodes attempt to communicate concurrently the number of failed connections increased due to collisions. They demonstrated that this failure percentage can drop if STAs gather beamforming information by listening to the channel during other nodes A-BFT slots. The network throughput gains are only noticeable for a large number of concurrent flows but this is certainly an interesting proposal which could be highly suited for certain usage models of 802.11ad in an IBSS (i.e. ad-hoc P2P mesh). Definitely there are practical limitations however since STAs may not be able to sense each others sector sweep due to the hidden node problem, or that obstructions may created unpredictable channel conditions such that the information gained by intelligent listening is of pessimal use.

Performance Analysis of 802.11ad

Researchers from Intel in [37] present a comprehensive set of simulations of 802.11ad performance in different realistic environments and configurations such as LOS range test and environmental testing (living room and conference room). The range tests were used to validate link budget parameters. The living room tests were used to simulate a media center application with a 60 GHz capable set top box. Lastly, the conference room was used to simulate a meeting scenario with projectors, client stations such as laptops and handhelds each requesting access to the 60 GHz network. In their range test, they show that use of beamforming can overcome the extra path loss at 60 GHz to provide an average goodput of 1.612 Gbps at a range of 10 meters. In the living room test, they simulated an average goodput of 3 Gbps with delay ranging from 74 us to 6.4 ms with no packet loss. In the conference room tests, they simulated ten concurrent links with an aggregated throughput around 754 Mbps. Their simulations demonstrate that 802.11ad is poised to

see significant growth to provide multi-Gbps rates by leveraging the WiFi's existing presence in the marketplace. Overall the research presents a good overview of 802.11ad and its potential but ultimately the PHY rates were selected arbitrarily, and their simulations cannot account for the dynamic effects present in real world deployments. Their environment and range testing was used as the basis for some of the real world characterization done in this research.

Chapter 3

Theoretical Analysis

In this section mathematical models are created for 802.11ad performance by applying antenna theory and link budget analysis to calculate SNR at various distances. Once the SNR is calculated achievable MCS can be determined by referencing SNR requirements for WiGig single carrier MCS rates, shown in Table 3.3. This is used to establish a performance baseline for RvR which can be then compared with empirical data. Since the test platforms only support the control and single-carry modulation schemes calculations are done for MCS 0-12 and omitted for OFDM and LPSCPHY for brevity.

3.1 PHY Rate Calculations

Given the information shown in Table 3.1 it is possible to calculate the PHY rate of a given MCS using Equation 3.1.

$$TP_{PHY} = R_{sym} \times N_{CBPS} \times {N_{repetitions}}^{-1} \times R_{code} \times O_{blocking} \tag{3.1}$$

where:

TP_{PHY} = PHY data-rate (bits per second)
R_{sym} = Symbol rate (1760 $\frac{sym}{s}$)
N_{CBPS} = Number of bits per symbol
$N_{repetitions}$ = Number of redundant symbol repetitions
R_{code} = Coding rate
$O_{blocking}$ = Blocking overhead

Table 3.1: Calculated PHY Rate

MCS	Bit/Sym	$N_{repetitions}$	R_{code}	$O_{blocking}$	TP_{PHY}
0	1	32	0.50	1.00	27.50
1	1	2	0.50	0.88	385.00
2	1	1	0.50	0.88	770.00
3	1	1	0.63	0.88	962.50
4	1	1	0.75	0.88	1155.00
5	1	1	0.81	0.88	1251.25
6	2	1	0.50	0.88	1540.00
7	2	1	0.63	0.88	1925.00
8	2	1	0.75	0.88	2310.00
9	2	1	0.81	0.88	2502.50
10	4	1	0.50	0.88	3080.00
11	4	1	0.63	0.88	3850.00
12	4	1	0.75	0.88	4620.00

The symbol rate used for single-carrier rates is fixed at 1760 Symbols/second. The number of bits per symbol is dependent on the modulation scheme (e.g. 16 QAM results in 4-bits of information after demodulation). MCS 0 and 1 use symbol repetitions rates of 32 and 2 respectively to ensure robust and reliable reception before link training and refinement. LDPC results in some overhead to allow error detection/correction, so the coding rate is defined as the ratio of $\frac{N_{databits}}{N_{total}}$. Lastly, there is some blocking overhead to account for IFS which is fixed at 1 for MCS 0, and 0.875 for MCS1-12. The values of variables in Equation 3.1 are summarized along with calculated PHY rates in Table 3.1.

3.2 Path Loss Calculation

Using the wave theory of light, wavelength λ is defined as the ratio of the speed of light and frequency[38]:

$$\lambda = \frac{c}{f} \tag{3.2}$$

where:

c = the speed of light in a vacuum (3x10^8 m/s)
f = the frequency of EM radiation (Hz)
λ = the wavelength (m)

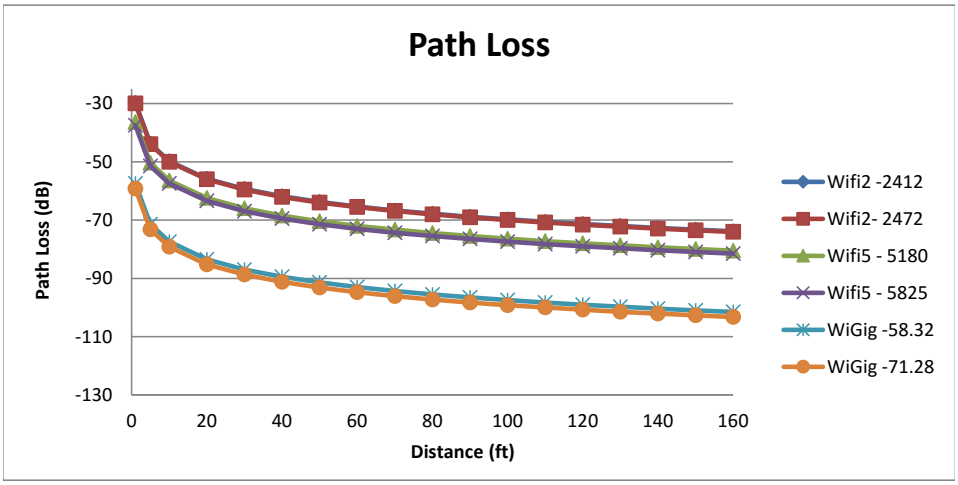

Figure 3.1: Free Space Path Loss Comparison (Wi-Fi vs WiGig)

In free-space electromagnetic signal propagation is governed by Friis transmission equation[39]. Friis equation, shown below in Equation 3.3, gives the received power by one antenna under idealized conditions given another antenna some distance away transmitting a known amount of power.

$$\frac{P_r}{P_t} = G_t G_r (\frac{\lambda}{4\pi R})^2 \tag{3.3}$$

where:

P_r and P_t = the receive and transmit power (W)
G_r and G_t = receive and transmit antenna gains (wrt an isotropic radiator)
λ = the wavelength (m)
R = the distance between antennas (m)

The inverse of the third factor is known as the free-space path loss (FSPL). The FSPL is plotted in Figure 3.1 for several 2.4/5 GHz Wi-Fi channels, WiGig channel 0 ($f_c = 58.32GHz$) as well as one centered at $f_c = 71.28GHz$ to investigate FSPL differences between spectrum allocation for the unlicensed mm-wave band, and the licensed 71-76 GHz band. Path loss is $\propto$ f hence between 5.825 GHz and 58.32 GHz there is approximately 20 dB difference in just the FSPL which is consistent with what is presented in [8, 40, 41].

Equation 3.3 can be alternatively represented on a log-scale with power in dB and gain in dBm as follows:

$$P_r = P_t + G_t + G_r + 20\log(\frac{\lambda}{4\pi R}) \tag{3.4}$$

These simple forms apply under the following ideal conditions:

- R >> λ
- unobstructed free space (no multipath)
- P_r and P_t are at the terminals (after any cable loss)
- P_t and P_r will only be fully delivered to/from the transmission medium (freespace) if the antenna and transmission line are conjugate impedance matched
- Antennas are correctly aligned and have the same polarization, the bandwidth is narrow enough that a single value for the wavelength can be assumed

The power delivered to the medium from the transmitter can be calculated as:

$$P_t(dB) = G_t(dB) + P_{T_{1x}} + 10\log(N_t) + L_{BF} \tag{3.5}$$

where:

$P_{T_{1x}}$ = Transmit Power per Antenna (dBm)
N_t = Number of Transmit Antennas
L_{BF} = Beamforming Loss (dB)

The basic form of Friis equation can be modified to account for non-ideal conditions, such as multi-path, impedance mismatch, differences in antenna orientation and polarization, and material absorption[42]. The modified equation is shown below in Equation 3.6.

$$\frac{P_r}{P_t} = G_t(\theta_t, \phi_t)G_r(\theta_r, \phi_r)(\frac{\lambda}{4\pi R})^2(1 - |\Gamma_t|^2)(1 - |\Gamma_r|^2)|a_t \cdot a_r^*|e^{-\alpha R} \tag{3.6}$$

where:

$G_{t|r}(\theta_{t|r}, \phi_{t|r})$ = is the apparent gain accounting for mismatch in antenna orientation
$\Gamma_{t|r}$ = antenna reflection coefficients for impedance mismatch
$a_{t|r}$ = polarization vectors to model for antenna polarization loss
α = absorption coefficient of the intervening medium ($\alpha_{freespace}$=0)

Friis equation gives the achievable signal power level seen by the receiver but a more important metric in wireless communication is the SNR at the analog-to-digital converter of the receiver. The SNR determines the relative signal level that can be extracted from a transmission given various sources of noise and interference.

3.3 Noise Floor Calculation

One source of noise is the thermal noise power P_{TN}, also known as Johnson-Nyquist noise[43]. This noise is generated by the thermal agitation of charge carriers inside electronics and conductors. When observing the output of spectrum analyzer there is a non-zero power-spectral density at 'unused' frequencies due to this noise. It can be modeled as additive 'white' gaussian noise (i.e. evenly distributed in frequency and amplitude). P_{TN} is present at all frequencies a larger channel bandwidth means more total noise power enters the receiver, hence $P_{TN} \propto BW$.

$$\begin{aligned} P_{TN} &= 10\log(kTBW) \\ &= 10\log(kT \times 1 \times 10^3) + 10\log(BW) \\ &= -174dBm/Hz + 10\log(BW) \end{aligned} \tag{3.7}$$

where:

P_{TN} = Thermal Noise Power (dBm)
k = Boltzmann's constant ($1.38x10^{-23}\frac{J}{K}$)
T = Temperature (Earth Surface $T_{AVG} = 290$ K)
BW = Detection bandwidth (Hz)

The total amount of noise received from interference, thermal agitation, and electronic components is known as the Noise Floor[43]. In the absence of interfering transmissions the Noise Floor can be calculated as:

$$\begin{aligned} NoiseFloor_{dBm} &= P_{TN} + NF_r \\ &= -174dBm/Hz + 10\log(BW) + NF_r \end{aligned} \tag{3.8}$$

where:

NF_r = The receiver noise figure

The receiver noise figure NF_r accounts for signal degradation caused by RF components in the radio signal chain.

3.4 Link Budget Calculation

Finally, the SNR can be calculated as the ratio of signal power to noise power[43].

$$\begin{aligned} SNR &= \frac{P_r}{P_{noise}} \\ SNR_{dB} &= 10\log(P_r) - 10\log(P_{noise}) \\ &= P_t + G_t + G_r + 20\log(\frac{\lambda}{4\pi R}) - P_{NF1} - P_{NF2} \end{aligned} \tag{3.9}$$

where:

P_{noise} = Total Noise Power (W)
P_{NF1} = Noise Floor Power (dBm)
P_{NF2} = Noise Figure (dB)

Table 3.2: RF Antenna Parameters

Parameter	Channel 0	Licensed Channel
F_c (GHz)	58.32	71.28
n	2	2
G_t(dB)	4	4
$P_{T_{1x}}$ per Antenna (dBm)	4.5	4.5
N_{T_x}	8	8
N_{R_x}	10	10
L_{BF} (dB)	0	0
NF_r (dB)	7.5	7.5
BW (MHz)	2160	2160

Figure 3.3 shows the calculated SNR vs Distance for both 60+ GHz channels. There is a slight decrease in SNR for the 71.28 GHz channel due to it having higher FSPL. It is important to consider that for these calculations medium absorption was neglected ($\alpha = 0$). Realistically in the mm-wave band the absorption factor may have substantial variations depending on the channel center frequency. The primary absorption molecules for mm-waves are H_2O, O_2, CO_2 and O_3. As shown in Figure 3.2 at 60 GHz gaseous O_2 has an absorption peak of 16 dB/km at sea level or 5 dB/km at 4 km altitude due to lower atmospheric concentration. In short range indoor LOS scenarios the effects are not

Table 3.3: Single Carrier MCS SNR Requirements

MCS Index	SNR(dB)	PHY Rate(Mbps)
0	-12	0
1	-0.35	385
2	0.71	770
3	2	962.5
4	3.3	1155
5	4	1251
6	3.64	1540
7	4.9	1920
8	6.35	2310
9	7	2502
10	11.1	3080
11	12	3850
12	13.3	4620

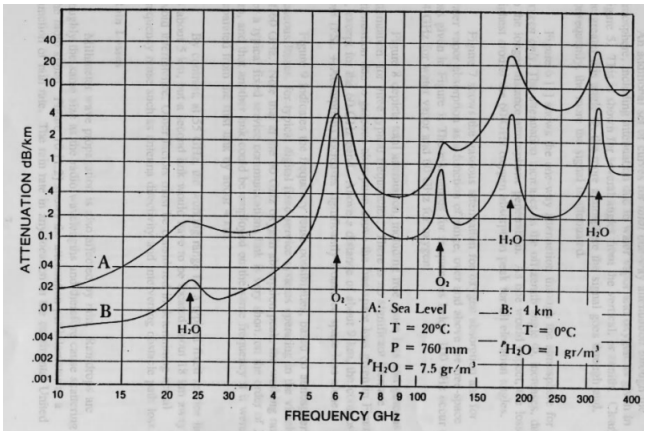

Figure 3.2: Atmospheric Absorption vs Frequency (Source [44])

as significant. The degradation is prevalent over longer distances and in outdoor conditions where the concentration is higher[8]. Furthermore, the effect of precipitation attenuation, or rain fade, during extremely heavy rainfall can contribute attenuation of approximately 5 dB/km which is only a small percentage of the aggregate attenuation in the 60 GHz oxygen absorption region [45]. Finally, Figure 3.4 shows the PHY rate vs Range using the SNR calculations based on the RF parameters given in Table 3.2. These are later used as a baseline to compare with empirical results.

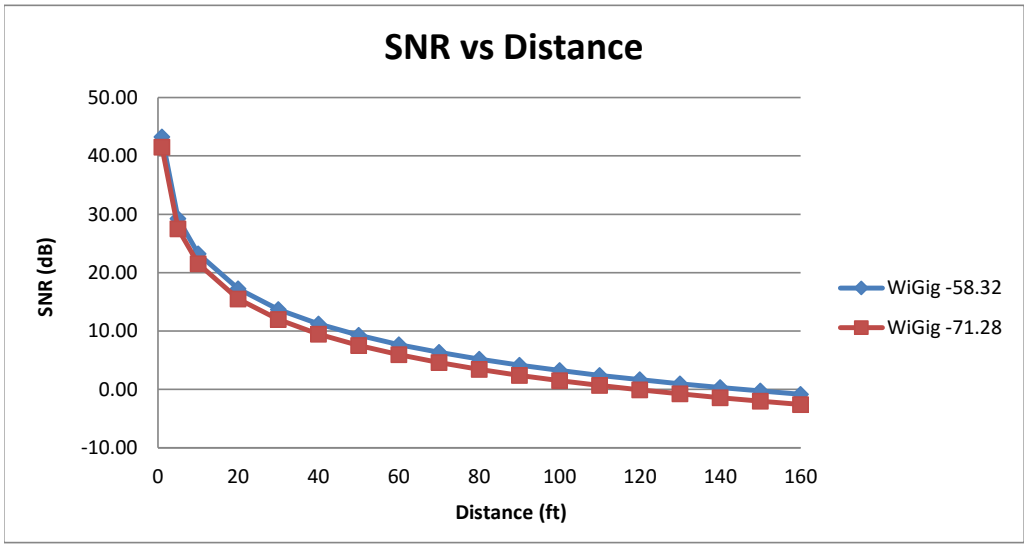

Figure 3.3: WiGig SNR vs Distance

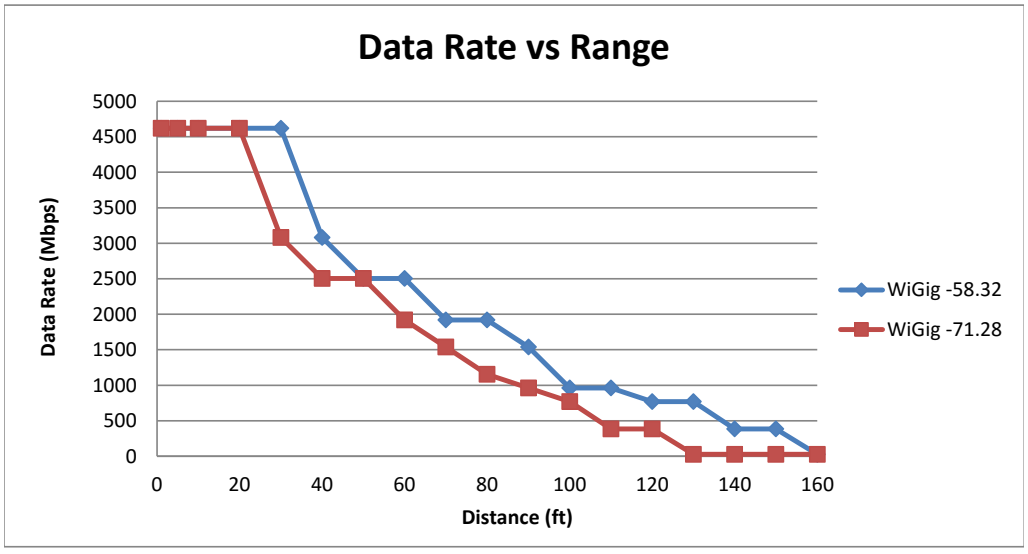

Figure 3.4: WiGig PHY Rate vs Distance

Chapter 4

Experimental Setup

This chapter outlines the test setups used for 802.11ad evaluation and characterization. This includes a variety of hardware such as reference design APs and early generation root devices, 11ad-enabled client laptop stations, wireless dock and tablet combos, and lastly the PoC 60 GHz platform which employed various antenna styles, placements and orientations.

4.1 Hardware Platforms

4.1.1 Antenna Topologies

In total there were two patch antenna topologies available for testing shown in Figure 4.1. The phased antenna arrays differed in the total number of antennas, the concurrently active transmit chains, and the use of a single antenna polarization, or a combination of multiple polarizations [i.e. cross polarized radio (XPR)].

During the first summer of testing there was one antenna topology available. This antenna (Falcon), pictured in Figure 4.1a, was an XPR which featured a combination of horizontally and vertically polarized antenna elements. The phased array contained 24 antenna elements but would only activate up to 8 simultaneously for increased beamforming gain. The use of XPR causes the service area to exhibit slightly stronger side lobes to supply a more omni-directional coverage pattern while facilitating improved performance and range along the bore-sight direction. The active power consumption of this topology was limited to 300 mW. The relatively low power consumption, along with the versatile

service area make this topology ideal for mobile stations such as portable battery powered devices, hence it is referred to as a 'client-style' antenna.

During the second summer of testing an additional antenna topology was available. The antenna was manufactured by a third-party supplier (SEMCO) and is pictured in Figure 4.1b. The SEMCO antenna contained only horizontally polarized antenna elements which provided very high directional gain at the cost of a steeper cut-off at high angles. Due to this the SEMCO topology was intended to be used in a diversity array, where multiple redundant antenna arrays were directed at different angles. With some overhead the AP would determine which antenna module was optimally directed for specific clients allowing it to leverage high directionality and mitigate the reduced coverage pattern of each individual element. During testing phase the baseband and PHY chipset only supported a single antenna module so testing along the strong coverage area was emphasized for this particular topology. There were several versions of the firmware which supported 250 mW, 500 mW and 1 W of active power consumption with varying number of active transmit chains. The phased array contained 32 antenna elements and if the high power firmware was used it could activate up to 16 concurrent chains. Due to the higher power consumption and high directional gain this antenna style is referred to as an 'AP-style' antenna.

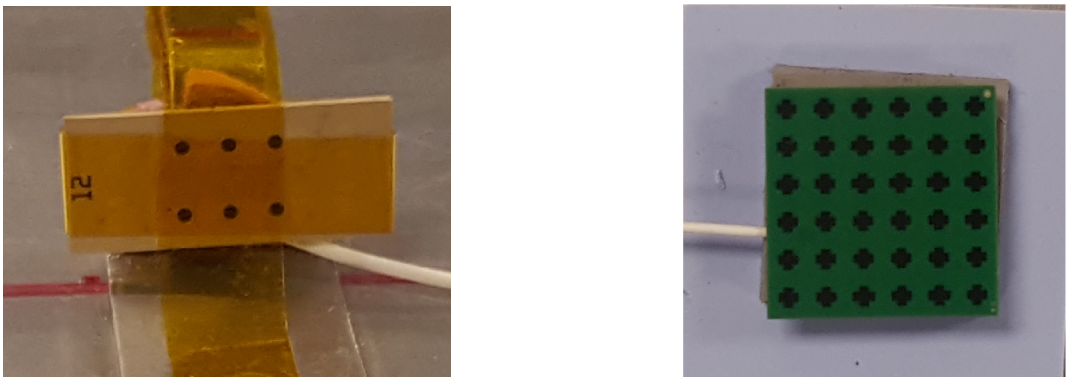

(a) Client Antenna (Falcon) (b) AP Antenna (SEMCO)

Figure 4.1: Antenna Topologies

4.1.2 Qualcomm/Atheros

Figure 4.2 shows two different form-factors of the tri-band reference APs featuring embedded ARM microprocessors. Figure 4.2a shows the exposed system with a detachable cable connecting to the isolated RF/antenna chip (shown in the middle mounted on the thin

gold plate). Along the right and left are additional cables for the 2.4/5 GHz antennas. Figure 4.3 shows an alternative reference design which supports only 60 GHz WiGig. The

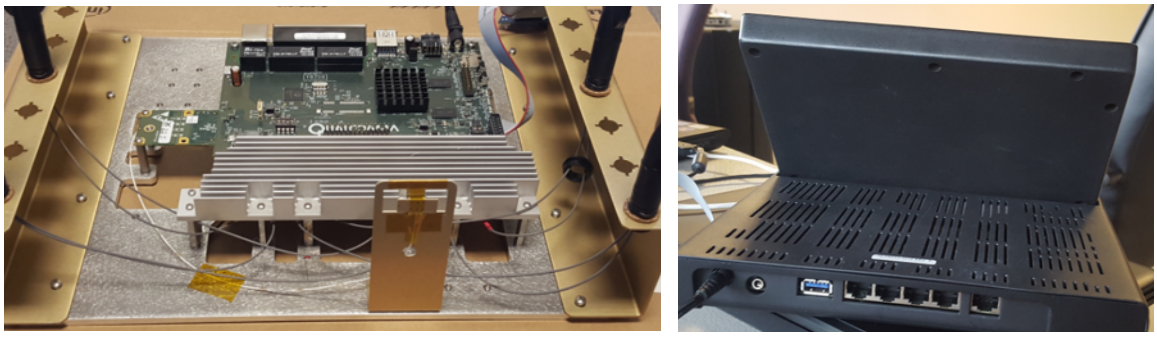

(a) QCA Tri-Band Reference Design (b) Alternative Form Factor

Figure 4.2: QCA Tri-Band Reference APs

host system is a Linux based system with a power Intel i7 processor. This minicomputer was used primarily in the first wave of testing before a custom platform was engineered.

(a) AP Housing (b) Exposed Antenna Module

Figure 4.3: Intel Minicomputer with QCA Falcon RF

4.1.3 Intel

Figure 4.4 shows a 60 GHz docking implementation from Intel. The dock and tablet pair using a simple GUI application in the Windows based tablet. Each of the interfaces present

on the PAN in Figure 4.4c would then appear on the tablet as if the dock acts as a wireless bus extension (WBE).

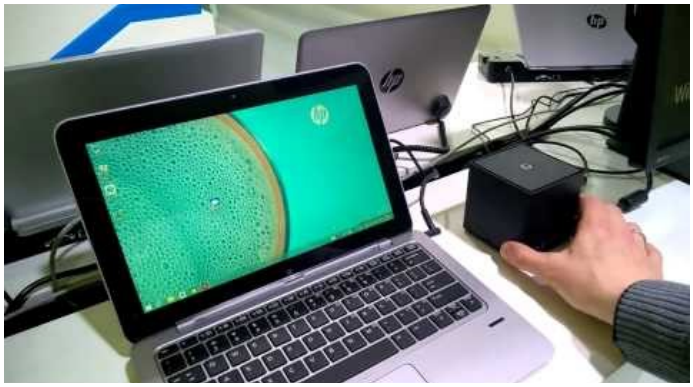

(a) Intel 60 GHz Dock and Tablet Combo

(b) PAN Interfaces

(c) 11ad Dock Usage Example

Figure 4.4: Intel 60 GHz Docking Solution

4.1.4 Client STAs

The WiGig enabled client stations are pictured in Figure 5.8. The Linux-based 6430u shown in Figure 4.5a was used during the first summer of testing (2015) during early characterizations. During the second wave of testing (2016) the three Linux-based E5250s (Figure 4.5b) were used in most tests as they displayed the highest performance. The Windows-based Travelmates in Figure 4.5c displayed bottlenecks at 2 Gbps so were used only in the multiclient tests as many stations were contending so the bottleneck was not apparent. In the Long Hall testing shown in Section 5.9 a set of testing was conducted with an E5250 as the client STA, and a set of testing was conducted with a second Cisco 3800 configured as a client STA to discover any limitations or bottlenecks with using the 3800 as a wireless bridge.

(a) Dell Latitude 6430u (b) Dell Latitude E5250 (x3) (c) Acer Travelmate (x3)

Figure 4.5: Client Laptops - Red Indicates Antenna Placement

4.1.5 Cisco Extension Module

A significant amount of effort was spent on characterizing the efficacy of numerous antenna topologies. Once a custom board was created which integrated Qualcomm's baseband and PHY chipset into a Linux-based AP operating system it was possible to examine the effects of several antenna styles with minimal design efforts since the antenna front-end attached to the PHY chipset via a detachable cable interface. In addition to testing two different antenna arrays, creation of a custom module enclosure allowed testing the effects of various front end design choices. Since 60 GHz waves have quasi-optical propagation and the EM interface to the AP was a side-car module, over the air (OTA) testing was conducted to determine if there were any AP-shadowing effects on the opposite side that the EM was mounted. A 'periscoped' antenna, shown in Figure 4.8c, was also deployed to determine if any shadowing effects could be mitigated by allowing the 60 GHz antenna module to protrude downward from the AP. Furthermore, several modules were created where the patch antenna was parallel to the ground plane, or alternatively perpendicular to the ground plane to investigate the affect of antenna orientation on coverage area. A comprehensive diagram of tested antenna styles is shown in Figure 4.8.

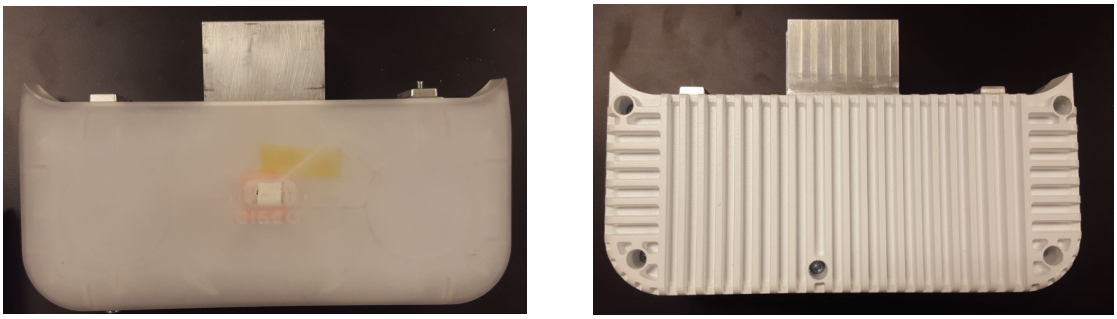

(a) Top View (b) Bottom View

Figure 4.6: 802.11ad Extension Module

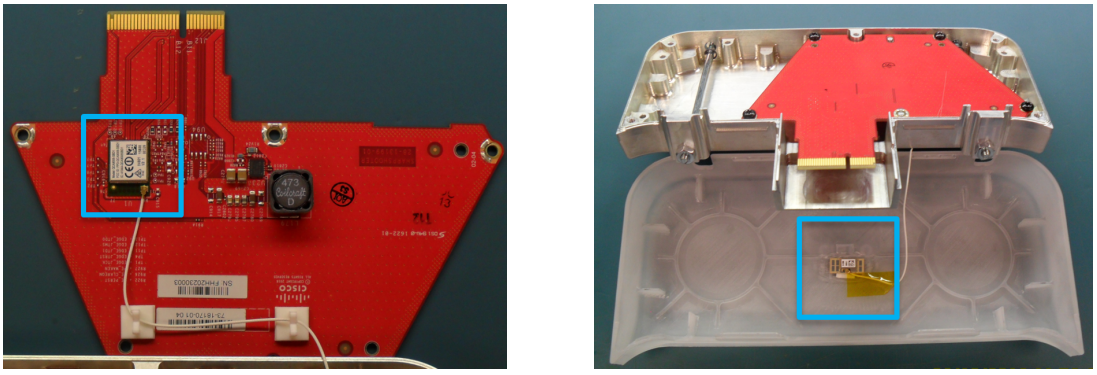

(a) Baseband IC (b) RF Antenna IC

Figure 4.7: Extension Module PCB View

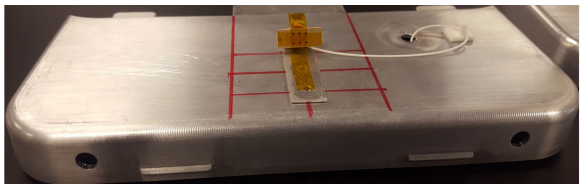

(a) Falcon - Perpendicular

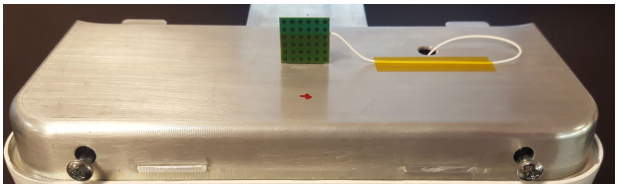

(b) Directional - Perpendicular

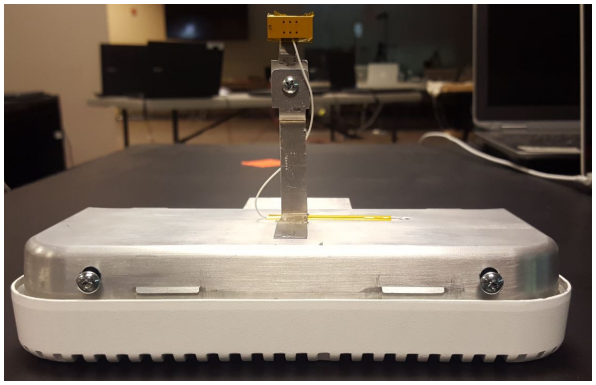

(c) Falcon - Periscoped Perpendicular

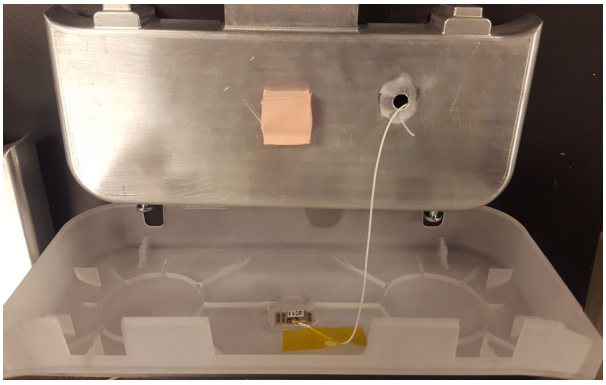

(d) Falcon - Flush (Detached Top)

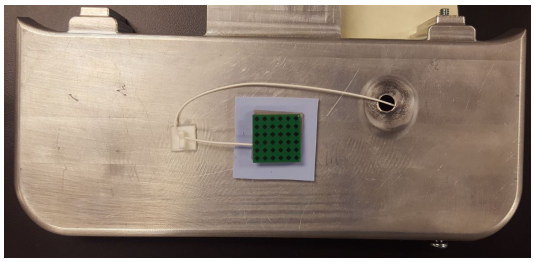

(e) Directional - Flush

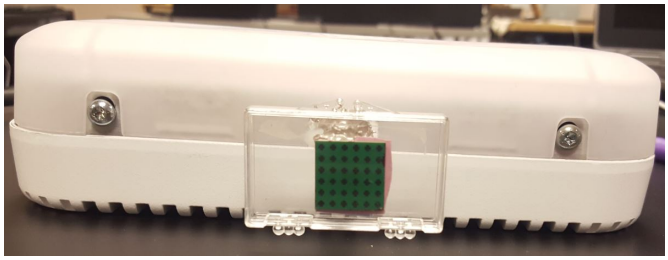

(f) Directional - External Perpendicular

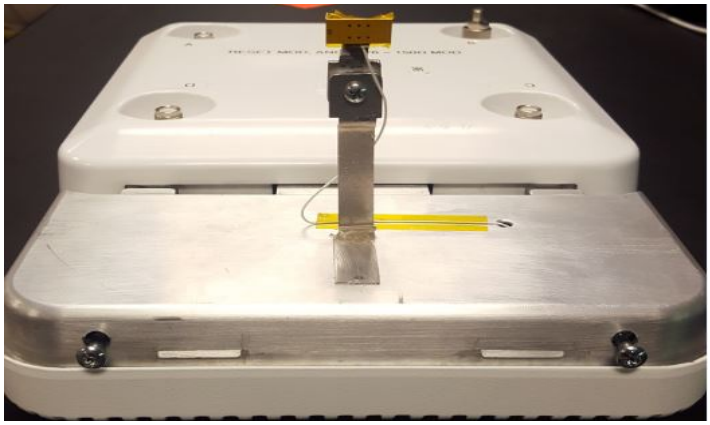

(g) Example AP with EM Inserted

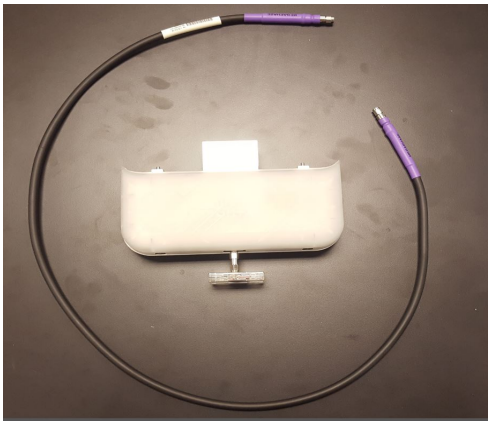

(h) Optional Extension Cable

Figure 4.8: Cisco Extension Module Antenna Styles

4.2 Test Environments

4.2.1 Meeting Room

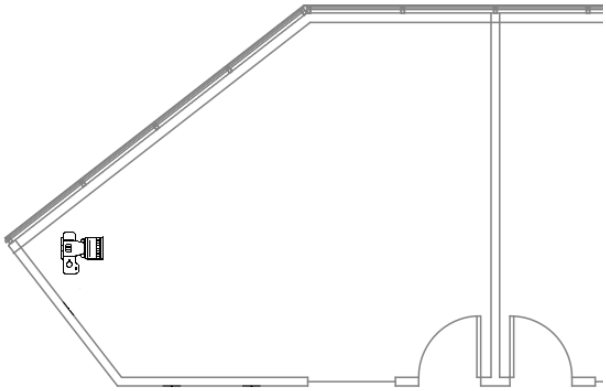

(a) Meeting Room Blueprint

(b) Alternate Meeting Room View

Figure 4.9: Meeting Room Environment

4.2.2 Open Office

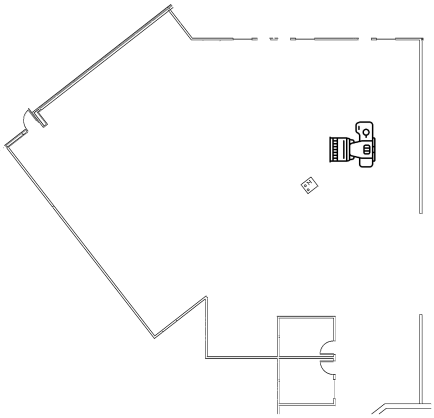

(a) Open Office Blueprint

(b) Alternate Open Office View

Figure 4.10: Open Office Environment

4.2.3 Cubical Farm

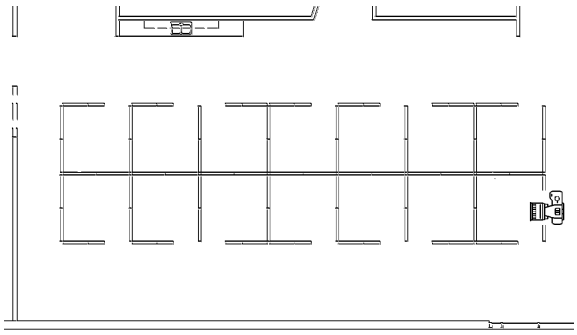

(a) Cubical Farm Blueprint

(b) Alternate View

Figure 4.11: Cubical Farm Environment

4.2.4 Long Hall

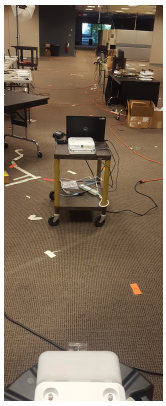

(a) AP View (8 ft)

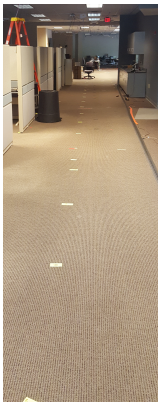

(b) STA View (130 ft)

(c) Side-View (8 ft)

Figure 4.12: Long Hall Bridge Setup

4.3 Test Methodology

The primary interest of the testing was to determine the capabilities of the technology, therefore the emphasis was on performance upper limits. There was a non-negligible amount of variation between test runs even with identical test configurations. For this reason typical tests comprised of running three test runs and selecting the best performing run as the one to characterize that point, environment, or antenna style. To eliminate errors due to beamforming convergence issues a brief dummy test was run once associated to ensure the beam alignment was optimized.

True Run Distance

The distances presented in test results are the calculated 'true run' distances to account for differences in endpoint elevation. In generic form:

$$D_{TR} = \sqrt{(x_{AP} - x_{STA})^2 + (y_{AP} - y_{STA})^2 + (z_{AP} - z_{STA})^2} \tag{4.1}$$

where:

D_{TR} = True run distance (ft)
(x_n, y_n, z_n) = Cartesian co-ordinates of AP|STA endpoints

In almost all the tests (excluding long hall and multi-AP), the AP was statically mounted on the ceiling (approximately 10 ft high). Similarly, in all tests except housing, multiclient and multi-AP, the client station was placed on either a rolling cart, upside-down garbage bin, or desk (approximately 2.5 ft high). The AP was considered the 'origin', hence $(x_{AP}, y_{AP}, z_{AP}) = (0, 0, 0)$, and Equation 4.1 was simplified to:

$$D_{TR} = \sqrt{{x_{STA}}^2 + {y_{STA}}^2 + {z_{STA}}^2} \tag{4.2}$$

Where z_{STA} is equal to 10 ft for housing and multiclient tests, 0 ft for multi-AP and long hall tests, and 7.5 ft for environmental and open office antenna characterization tests.

4.3.1 Software Tools

Connectivity Establishment

For establishing a wireless connection (scan, association and authentication) since WiGig is conformant to the 802.1x standard for network access control, connecting to the wireless could be done using the built in network manager in Windows or Linux. In Linux,

the authenticator endpoints were hostapd and wpasupplicant instances for AP and STA respectively. This meant in the absence of a network-manager (on some minimalist embedded systems such as the PoC platform) scan and association could be done simply using the command-line linux wireless tool 'iw'. Once associated for Layer-3 connectivity both endpoints had to be put on the same subnet using either static or dynamically assigned IP addresses. Layer-2 connectivity could be verified using the address resolution protocol tool 'arpping' then Layer-3 connectivity could be verified using the internet control messaging protocol tool 'ping', or a Layer-4 throughput measurement tool such as Iperf3.

Application Layer Throughput Measurement

Several different software tools were used to measure wireless bandwidth, performance and reliability. Netcat is a simple Unix utility which reads and writes data across network connections using TCP or UDP. It can be used to create a listening endpoint to which a sender can connect to and transfer a file, message, or random data from various input streams. Netcat is installed by default on many of the systems used, so at a rough estimate of wireless performance was established by generating large ($\sim$10 GB) files and transferring them across the network with netcat. Several other third-party networking tools were used to fine tune parameters such as test interval, packet size, total transfer size, TCP congestion window size. Other softwares that were tested include nuttcp, netperf and Iperf/Iperf2/Iperf3. In the end the results were generally consistent across all the softwares used, and it was discovered that wireless throughput was not substantially impacted by modifying window/packet sizes in these single hop network configurations. Due to Iperf3 being open-source, free, and actively maintained it quickly became the paragon for network performance profiling. Iperf3 contains a wide array of features inherited from the previous generations of Iperf, along with important features that were present in other tools such as nuttcp and netperf.

Test Runs

Once connectivity was established, wireless performance was tested by setting up a listening server, and then getting the other endpoint to act as a client to connect to the listener. Both uplink and downlink tests could be initiated from the client by including or omitting command-line switch -R for reverse mode. Whether or not the AP or client was setup as the listening server depended on the test suite being run. For example, for multiclient tests where there are many clients, it was desirable to have the clients setup as listeners so that simultaneous connections could be quickly initiated from the central AP by starting

sending processes and then pushing them to the background with standard shell syntax (e.g. directing output to a logfile and including the and directive to create a background process). The alternative would be to setup the central AP listening to a range of ports, and then to go to individual clients and setup senders to their respective listener ports. For clients that were spatially separated this created a delay in the time which the connections start which corrupt the start of some of the tests until all links are active. Iperf3 has a switch (-O N) to omit the first N seconds of a test from the results but ultimately it is much more prone to error than simply initiating the tests from the central AP. It takes a negligible fraction of a second for a modern system to fork a newly created process to the background. For the spatial re-use testing (multi-AP) it is not possible to initiate all tests from one centralized point since there are two independent (i.e. not interconnected) links. For the multi-AP tests, the first test would need to be initiated then the tester would have to travel to the second endpoint to initiate the second test. To mitigate the travel time, the first test was initiated with a compound statement such as 'sleep 4 && iperf3 -c 10.0.0.1 -i1 -t30 -u -b0' which would allow the tester four seconds to walk from the first STA to the second STA so both tests could start approximately coincidentally. Lastly, there were times where it was favorable to have the AP act as the listening server such as tests where the AP was statically mounted on the ceiling and the client would need to be moved around (either during the test, or between test points). In these cases, it was ideal to be able to re-locate the client, and run a short dummy test to ensure beamforming configuration was optimal before a profiling measurement was made.

Additional Statistics

In addition to being interested in statistics output by Iperf (primarily throughput), the radio driver used for the 11ad chipset includes standard Linux debugfs files for more link-level characterizations. It is possible to periodically poll these files to profile additional measurements like MCS rates, signal quality indication (SQI), effective PHY data rate, and the drivers estimate of application level TP (Goodput). The driver reported Goodput was similar to that reported by Iperf3 but was consistently lower by roughly 50 Mbps. Since the Goodput was periodically polled, it contained some sampling error due to processor scheduling inaccuracy, whereas Iperf3 is aware of exactly how much data is successfully sends across the network due to deterministic packet sizes and numbering. For this reason the TP reported by Iperf3 was favored over the driver Goodput.

Aggregate Data Streams

There were scenarios where the driver Goodput was favored over Iperf3 throughput for simplicity of data processing. Due to processing power limitations on the 3800 the data stream was divided across both processor cores to mitigate the processing bottleneck. In the multiclient tests, this was done for up to 6 clients, creating a total of 12 Iperf streams. In this case, it was much more simple to use the driver reported Goodput in lieu of the aggregated Iperf3 TP and accept a small sampling error. Furthermore, before Iperf3 had been compiled for the 3800 it could not be used as an Iperf3 endpoint. To stress test the wireless link the 10 Gbps wired interface of the 3800 was connected to a Cisco 10GigE Catalyst switch, where the traffic was de-aggregated to four wired endpoints. Since the wireless bandwidth was split in four, the Goodput was favored to get a rough estimate of the cumulative performance. Eventually, Iperf3 source was compiled for the 3800 OS allowing it to act as an endpoint, but ultimately it still suffered from the aforementioned limitation in processing power.

Extensive Use of Scripting and Automation

Since several things were monitored during tests at the same time, various shell scripts were created for both local and remote initiation of statistic logging and traffic generation/termination. Due to testing being conducted on a variety of systems, the shells were not always the same. The PoC platform used a minimalist busybox/ash shell whereas the ubuntu client stations used Bash, so scripts needed to be slightly modified in order to work correctly. Once all the test data was collected (Iperf output, ping latency, link-level statistics), various scripts were used to parse and process the data into delimited values which could be imported into Excel for infographic generation. All of this automation ensured uniformity of testing, and enabled for tests to be initiated and left alone for significant amounts of time, where the scripts could detect and recover from test/link failures, launch advanced tests (such as multiclient), and save processing time for interpretation of results.

Chapter 5

Results

The first section of this chapter describes the system design engineering challenges associated with development of 802.11ad enabled hardware systems. This includes some of the difficulties associated with software integration, hardware limitations and bottlenecks, and difficulty of testing these early generation platforms. It then outlines a detailed test configuration and analysis of results for a wide array of test suites. Each subsection contains discussion of results along with qualitative and quantitative observations and explanations.

5.1 Limitations and Encountered Issues

5.1.1 Issues and Workarounds

There were many difficulties encountered when trying to engineer a custom PoC 802.11ad module for the Cisco 3800 AP. The first major hurdle was integrating Qualcomm's open source 802.11ad driver (wil6210) as a loadable kernel module in the AP's Linux-based operating system (OS). This was done by configuring the OS buildroot to package the driver, but the driver source code had to be modified due to discrepancies between the Linux kernel versions. The mainline driver is updated to the latest Linux kernel but commonly embedded systems are developed on slightly older but more hardened kernel version. To get the driver to compile it first had to be 'back-ported' to the correct kernel version. After this, the driver had some dependencies which were not packaged into the OS already. This required addition of the Linux networking library netlink, regulatory support for 60 GHz, and some user-space configuration tools such as iw, hostapd, and wpa_supplicant.

Once the driver was bundled as a kernel object and the user-space configuration tools were added, it was possible to configure the 3800 as an AP or STA.

The next difficulty was that the OS image did not include any benchmarking applications. Once clients were associated to the AP it was possible to verify connectivity using ping and address resolution protocol. In order to get an early estimate of wireless performance, tests were run between two wireless stations. Since the data was sent from the first client to the AP then from the AP to the second client, the number of wireless packets had to be doubled. Therefore the measured performance was approximately half of the actual single-hop performance. These early tests showed double-hop throughput around 1.4 Gbps, which was consistent with estimates. It was then verified that throughput over 2 Gbps could be obtained by using a single-hop configuration and passing a randomly generated file over netcat, since netcat was packaged in the OS image. This was also verified using wired endpoints, since the 3800 contains a 10 Gbps Ethernet interface. There are very few devices which support over Gbps Ethernet, so the wired interface on the AP needed to be attached to a high performance Cisco Catalyst switch. Several Gbps wired endpoints were then attached to the switch on the same virtual LAN. Traffic tests were run between the wireless client and several wired endpoints. The aggregated throughput verified the wireless maxing out around 2.8 Gbps. In the end, Iperf was compiled for the 3800 OS by capturing the compilation flags for other linux packages. Once Iperf source was manually built for the host-OS the 3800 could be used as an endpoint in wireless tests.

Another issue encountered was related to bottlenecks on the AP. The wireless clients tested all had powerful Intel i7 processors. The AP used an embedded multi-core ARM processor which offers less processing power than an i7. It was also discovered that certain tests would perform better when running on a certain core within the CPU. This was due to the other background processes that were running on the AP for configuration and management of regular Wi-Fi functionality. To alleviate this bottleneck performance tests were split across both cores and the aggregate throughput was taken. The processing power difference also created asymmetrical uplink and downlink capabilities. On the downlink, the AP was unable to generate traffic as quickly as the client laptops. On the uplink, the AP was unable to process the traffic as quickly as the client laptops. This prevented full saturation of the wireless link-layer capabilities. It was discovered that the uplink bottleneck could be somewhat mitigated by having a wired endpoint in the test as well. In the uplink tests, two performance tests were run. One was from the STA to the AP, and the second was from the STA to a wired endpoint on the AP. This offloaded some of the traffic processing to the wired endpoint, resulting in more processing power on the AP used for link-layer acknowledgments. This minor enhancement allowed higher throughput on uplink tests which was ultimately why uplink was emphasized for the PoC platform. A

wired endpoint could not be added on the station, as you cannot bridge a wireless client for security reasons. To add a wired endpoint on the station it would have to be done using two non-overlapping subnets and layer 3 routing. However, since the routing is done at layer 3 (as opposed to a layer 2 bridge) there is one more layer of interaction with the networking stack, diminishing any benefit to offloading the packet processing to a wired endpoint. In the end, only the uplink bottlenecks could be reduced without making low level software enhancements to the networking subsystem.

There were also many difficulties associated with automation of test runs and parsing of test results. For example, the multiclient tests needed to be initiated simultaneously and the clients are spatially separated. In order to do these tests they were initiated from the wired endpoint on the AP but with the '-R' reverse flag to test the uplink. The scripts also needed to poll the config files for PHY rate, MCS, SQI etc. The scripts would invoke several background processes to run several tasks simultaneously. All of the tests were automated using shell scripting, where the shell depended on the platform (ash/busybox or bash). The use of advanced scripting techniques would allow detection and recovery from test failures as well. For example, if a test were to timeout after a certain period of time, or fail for some other reason, return codes from the commands could be used to re-initiate the same tests. Another example is the multi-AP tests. There was no connectivity between the links so tests had to be initiated from both stations simultaneously. Due to the travel time between station's taking X seconds, the first test run would be prefixed with sleep X, so that they would start at approximately the same time. The tests would produce a plethora of data. There were typically three runs taken for each test, and each run would produce an Iperf log, ping log, and configuration log. Multi-purposed helper scripts were used to initialize listening servers remotely, quickly summarize test results and to parse and extract the data-points for Excel processing. Lastly, for the open office antenna characterization testing, a MATLAB script was used to produce the topographical throughput maps. The script accepted a text file with x,y,z tuples where x,y were spatial coordinates representing the test locations, and z were the measured values for throughput, PHY rate, or MCS for example.

5.1.2 Software Bottlenecks

It is important to note that even leading edge systems are not optimized for multi-Gbps throughput and the VHT offered by 802.11ad can be bottlenecked by the host system. When receive queues fill at a high rate, it requires a lot of processing on the system side to avoid buffer overflow (for example, large memcpy in a Linux system). Similarly, as ingress packets generate message signaled interrupts (MSI) on the PCIe bus, interrupt handlers

are triggered very frequently causing a lot of context switching which can further reduce performance. Lastly, modern OS networking stacks are not optimized for higher rates. The Windows client laptops tested would only achieve throughput approaching 2 Gbps because the default Windows networking stack is not meant to handle voluminous data at such high rates. The default Linux networking stack also suffers from similar limitations. In specialized products, datapath optimizations must be made to accelerate some of the functionality of the networking stack. When the OS networking components are offloaded by either hardware (HardMAC) or software (SoftMAC) intellectual property the acceleration is offered by what is termed as a network subsystem (NSS). Hardware based acceleration can be accomplished by dedicated hardware on the network interface card (NIC) which receives commands via firmware and PCIe mailboxes. Software acceleration can be achieved by using various techniques such as more efficient memory allocation/deallocation, using larger queues/buffers, or using huge virtual pages for example. It is possible to do minor tweaks to the TCP/IP stack in Linux and Windows but they provide marginal benefits in terms of performance or reliability in real-world scenarios. Ultimately they were not a panacea and could not alleviate issues with multi-Gbps data processing on these systems.

5.1.3 Hardware Bottlenecks

Another limitation with most systems is host bus which connects to the 802.11ad baseband module. Modern systems use PCIe for high-performance adapters which connect to a systems front-side bus. Different generations of PCIe standard provide different throughput. The unidirectional throughput of Gen1, Gen2 and Gen3 are 2.5 Giga transfers per second (GT/s), 5 GT/s and 8 GT/s respectively. Actual data rates can be calculated by multiplying the raw transfer rate by the ratio of data bits to total bits, since overhead bits are encoded to ensure error-free transfer. The encodings used are 8b10b for Gen1 and Gen2, and 128b130b encoding for Gen3. The performance of a single lane PCIe interface (x1) is summarized in Figure 5.1 for PCIe Gen1, Gen2 and Gen3.

Conventional 802.11 Wi-Fi cards are considered low-speed peripherals and as such have not encountered practical limitations due to saturating a single PCIe Gen1 lane. This assumption is nullified by the massive data rates supported by 802.11ad. Most early iteration 802.11ad devices support up to MCS 12 which has a physical data rate of 4620 Mbps. Furthermore once OFDM rates are supported data rates increase up to 6756.75 Mbps @ MCS 24. These unprecedented rates impose new challenges in terms of system architecting and require special design considerations. As shown in Figure 5.1, PCIe Gen1 can be saturated by MCS7, Gen2 can be saturated by MCS11, and Gen3 can be saturated by MCS24. As early platforms supported up to MCS 12, the test results presented here are effected by

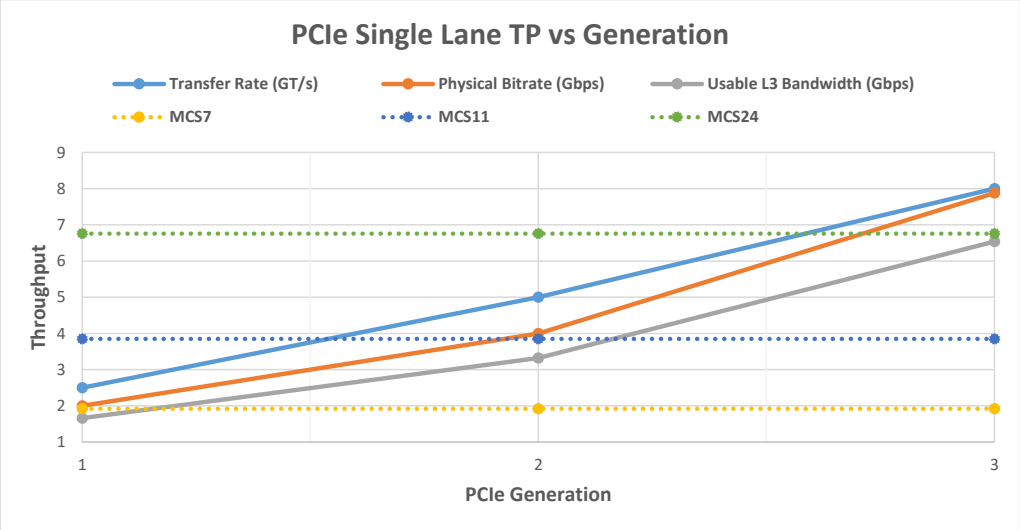

Figure 5.1: PCIe Single Lane Throughput vs Generation

the bottleneck considering the PoC platform deployed only a PCIe Gen2 link. Initially the PoC platform was linking the EM as Gen1, which resulted in observed bottlenecks in tests. This was quickly debugged by fixing some code in the uBoot bootloader.

5.2 Performance Benchmarking

The intent of this test was to establish an upper limit for performance and reliability in a close-proximity LOS environment. This is to establish the practical limit of the technology so that relative comparisons can be made with non-ideal scenarios. In this test, the AP and STA are co-located with antennas oriented so their maximum EIRP directions coincided. It is observed in Figure 5.2 that the application layer throughput reported by Iperf and the driver Goodput are very close. Iperf seems to display more distinct variations and has higher accuracy, due to sampling accuracy associated with polling the debug file which displays average goodput. Iperf is aware of exactly how much data is sent across the network due to uniform packet sizing and sequence numbering, whereas the Goodput likely outputs a sliding average calculated on the radio chipset - which is prone to sampling error due to minor variations in the interval between samples (due to scheduling, CPU load etc). It is also apparent that the Goodput figure reported is slightly higher than the Iperf throughput, by approximately 100 Mbps. This is likely due to Iperf subtracting

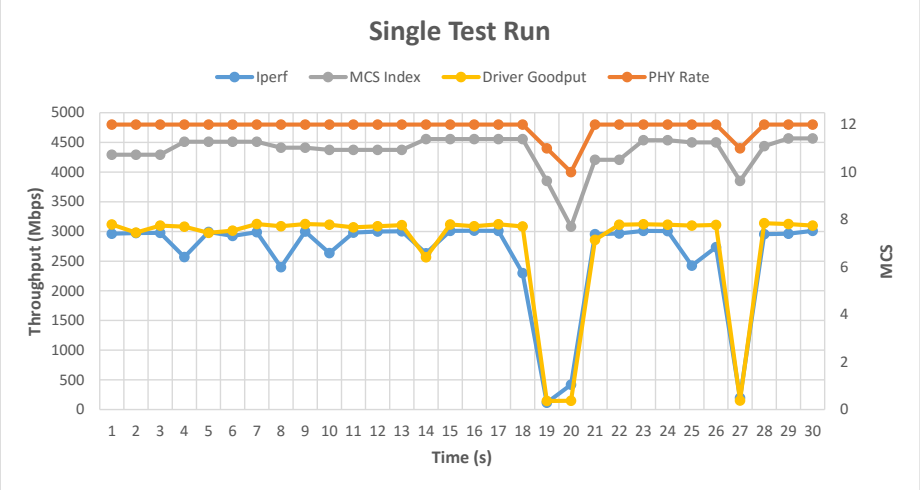

Figure 5.2: Close Proximity Test (d = 8 ft)

a small amount of overhead when computing the effective throughput. In Figure 5.2 it is apparent that even in close proximity static LOS testing there can be significant variation in test performance likely due to rate-shifting algorithms implemented on the target systems. Between 18-20s and at 27s, the data-rate drops substantially momentarily. Typically down-shifting the rate will occur when a certain threshold of failed receptions occur in a sequential window. Since the link quickly recovers it is probable that the link shouldn't have been drastically downgraded. It is also possible, in this close proximity environment, that the receiver is being over-driven. Receiver overdriving occurs when the input signal has such high amplitude that it saturates the low-noise amplification stage of the RF front end. This can lead to clipping and distortion of the signal inhibiting proper reception. It is likely that some overdriving caused a few unsuccessful transmissions, which caused the rate shifting algorithm to downshift to MCS 8. After several successful bursts at MCS 8 the rate-shifting algorithm determines it can return to a higher rate.

5.3 Intel Docking Station

The docking stations provided by Intel did not have full networking functionality available so could not be configured for networking. Instead, the Intel docks act as peer-to-peer (P2P) docking stations for an 802.11ad enabled tablet. Intel's primary investment with WiGig has been for docking applications where high bandwidth peripheral cables are at-

tached to a dock which wirelessly pairs client tablet/laptop. The ports present on the Intel dock show in Figure 4.4c include several USB 2.0 interfaces, VGA/DisplayPort/HDMI, audio, and gigabit ethernet. It was extremely convenient to be able to carry the laptop around the room while maintaining connection to a wireless display, external storage, and a wired network. Intel has correctly identified that WiGig has a significant application space eliminating wires in the personal area network (PAN). Since networking wasn't enabled yet on the docks, and the wired interface was simply a gigabit ethernet, throughput over 802.11ad was tested by using a simple file transfer to several USB 2.0 drives. At the time of testing, each drive would only support roughly 50 MB/s write speeds. To overcome this, a total of four flash drives were used with simultaneous file transfer and the Windows task manager was used to estimate the average write speed. In total, the 802.11ad dock implementation could saturate three of the Kingston USB 2.0 writes, and the fourth one would be limited from its maximum bandwidth. The total effective transfer rate was between 1.3-1.4 Gbps. It is possible that multiple USB interfaces are sharing the same bus or controller, so in hindsight a TCP transfer over the Gigabit ethernet should have been added to verify the performance limitation. Ultimately, due to the lack of controllability of performance testing on the docks testing of them was limited to this early characterization. In summary, the docks do an exceptional job at eliminating all wires from a portable PC with the exception of the power cable. It will be interesting to see advancements in the coming years for wireless display (WiDi) over 802.11ad and also processor specific enhancements Intel can add to fully leverage this high-performance technology.

5.4 Environmental Characterization

The aim of these tests were to identify suitable environments for 802.11ad deployments. Different network topologies create unique challenges and limitations for wireless; the intent of these tests were to measure WiGig throughput and reliability in various real-world environments.

5.4.1 Conference Room

The conference room is a small ~30 ft diameter room which is similar to a home entertainment center or a meeting room. The close proximity of the walls present multiple propagation paths from AP to STA so it is a good environment to characterize LOS performance under the effects of multi-path interference. A home entertainment center typically contains several bandwidth hungry appliances such as docking stations, smart TVs, media

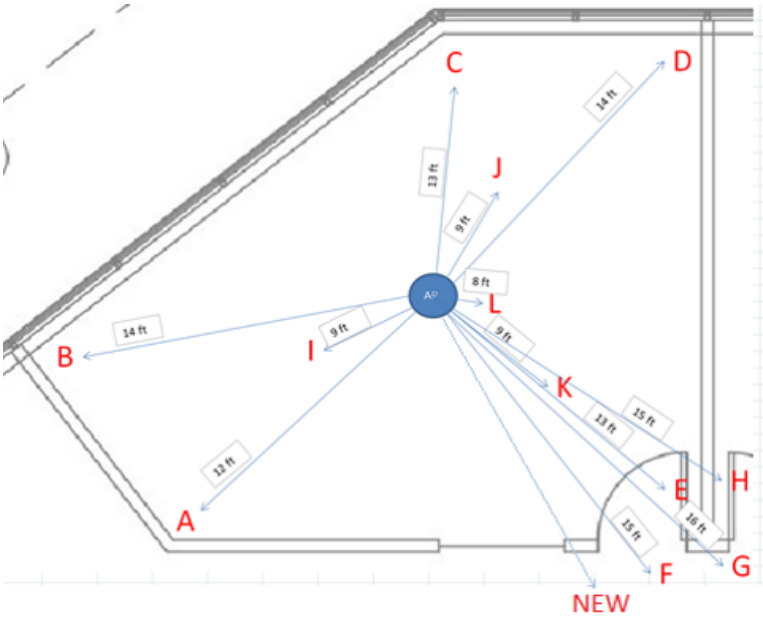

Figure 5.3: Conference Room Points

Table 5.1: Conference Room Performance (Distance in feet)

			Extension Module Performance	
TestPoint	LOS?	Distance	TP (Mbps)	Std Dev (Mbps)
A	Y	12	1886.83	35.20
B	Y	14	1806.63	229.72
C	Y	13	1904.17	17.87
D	Y	14	1865.87	82.14
E	Y	13	1936.91	69.06
F	Y	15	1902.94	14.53
G	N	16	0.00	0.00
NEW	N	16	580.19	204.64
H	N	15	0.00	0.00
I	Y	9	1904.11	10.04
J	Y	9	1906.26	9.25
K	Y	9	1901.44	20.76
L	Y	8	2350.58	41.34
	Average	12.54	1813.27	66.78

servers, laptops, and handheld devices like tablets and phones. In a typical teleconferencing application there are telephony appliances like microphones, video recorders, HD displays and numerous client devices like laptops and smartphones. The high bandwidth of video streaming requires codecs to compress A/V information to not saturate wireless links. This creates additional latency overhead which can further diminish effectiveness.

The results in Table 5.1 suggest that WiGig is highly suited for close-promixity LOS environments such as a meeting room. The average throughput was nearly 2 Gbps with very high reliability indicated by the low standard deviation. The high throughput and low latency demands of telephony and media streaming make the short range 60 GHz band the prime candidate. High frequency re-use allows WiGig act as a high performance small-cell network with minimal interference into adjacent rooms. The fact that test point 'NEW' could maintain approximately 600 Mbps performance by using a first-order reflection off the open door also suggests that the data-rate drops off in the presence of obstructions or blockages.

5.4.2 Open Office

The open office environment is similar to a dense-deployment like an office building where there are many devices in a large open space. The long distance from AP to the walls reduces the effects of multi-path propagation; which could potentially be detrimental for 802.11ad in the presence of obstructions such as humans or pillars.

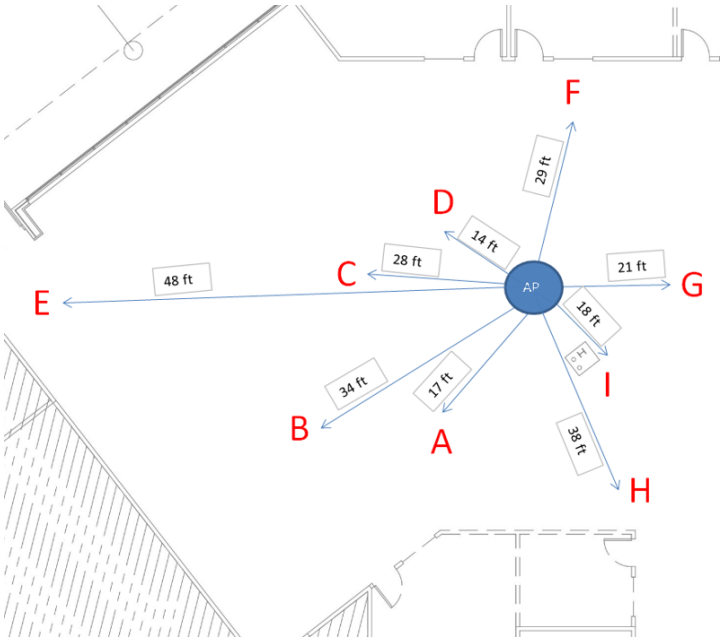

Figure 5.4: Open Office Performance

Table 5.2: Open Office Performance (Distance in feet)

			Extension Module Performance	
TestPoint	LOS?	Distance	TP (Mbps)	Std Dev (Mbps)
A	Y	17	1904.68	16.38
B	Y	34	1591.39	271.21
C	Y	28	1903.00	14.79
D	Y	14	2424.26	737.42
E	Y	48	1163.47	334.46
F	Y	29	1883.26	82.23
G	Y	21	1875.33	40.94
H	Y	38	884.92	142.43
I	N	18	1060.17	181.10
	Average	27.44	1632.27	202.33

The results in Table 5.2 indicate that WiGig is strongly suited for large open-office spaces. The absence of support columns, cubical dividers and other obstructions allows WiGig to maintain very high performance LOS links. The average throughput decreased slightly compared to that of the close-proximity conference room and the reliability decreased but overall performance and stability were still very high, especially considering the average endpoint separation was over 25 feet. As expected performance slightly decreased at high angles from the AP antenna bore-sight which suggests that the infrastructure devices should be strategically located such that their maximum range directions coincident with the longest dimension of the room or where the highest load stations will reside. The NLOS test point 'I' which was blocked by a support column, displayed very low performance since there are no close walls for the signal to reflect off of. Conversely, in the conference room the close wall boundaries allow nearby reflection surfaces to mitigate obstructions. This is atypical for Wi-Fi as multi-path is usually a caveat since it causes delayed versions of signals to arrive at the receiver which can destructively interfere. With WiGig multi-path may actually be beneficial to overcome obstacles while not causing significant degradation to LOS performance.

5.4.3 Cubical Farm

The cubical farm is also a dense-deployment scenario but presents unique challenges for 60 GHz waves due to obstruction from cubical walls, cubical orientation, humans and other objects. Most of the test points in the cubical farm are NLOS which creates a difficult environment to maintain high throughput with 802.11ad.

The results in Table 5.3 display promising performance and reliability of 60 GHz links in the cubical farm. On average, the cubical farm locations performed similarly to the LOS

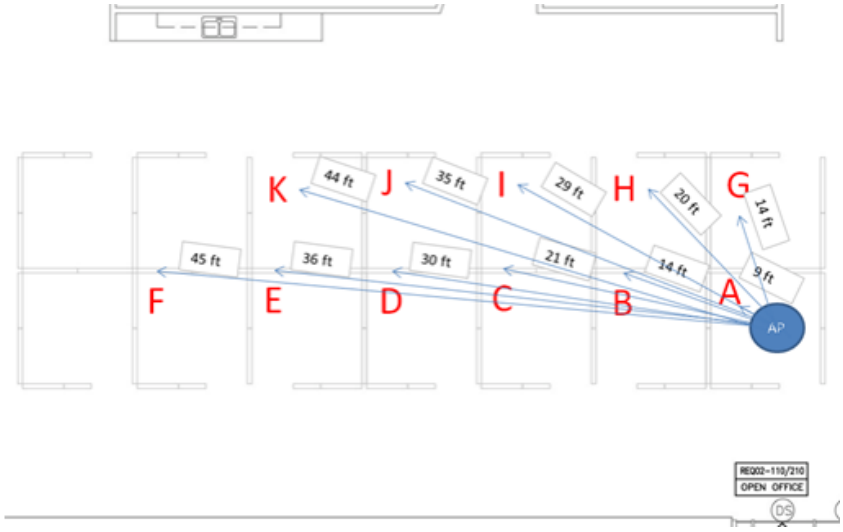

Figure 5.5: Cubical Farm Performance

Table 5.3: Cubical Farm Performance (Distance in feet)

			Extension Module Performance	
TestPoint	**LOS?**	**Distance**	**TP (Mbps)**	**Std Dev (Mbps)**
A	Y	9	2429.12	452.08
B	N	14	1387.01	193.26
C	N	21	1882.73	95.25
D	N	30	1852.98	101.40
E	N	36	1894.39	25.55
F	N	45	915.20	98.95
G	N	14	1859.27	37.21
H	N	20	1774.09	242.25
I	N	29	1529.14	94.61
J	N	35	1794.01	168.63
K	N	44	831.32	186.71
	Average	27.00	1649.93	154.17

open office space. This could potentially be due to the fact that all test locations were along the bore-sight direction of the AP antennas radiation pattern. Reliable multi-Gbps performance in this NLOS environment was not expected but some negative environmental effects were observed. Firstly, the client stations were placed within the cubical at a location which displayed the highest performance while not being LOS. This was in an effort to identify the capabilities, or upper limit of achievable performance in this NLOS environment. Substantial variation was noticed based on the clients inter-cube location: position, elevation, and orientation. Finding the highest performance placement was rel-

atively straightforward. By envisioning a first-order reflection off the cubical wall a user could identify which region of the cube elicit highest performance or which areas suffered anti-nodal shadowing. It was also noticed that as the station approached the boundary between N/LOS the performance would gradually increase until it became that of a bore-sight LOS test. This suggests that the 60 GHz links can be maintained and strengthened through the combined use of reflection and diffraction.

5.5 Housing Type

The aim of the housing tests were to determine signal degradation effects due to AP housing type. In these tests the AP and STA antennas were set at fixed orientation and distances, and the enclosure type was varied. Some APs and home routers use external antennas to isolate antenna front-ends from noisy components such as power supplies. Enclosing phased antenna arrays inside a casing also can cause reflections and refractions which reduce the antennas ability to beam-steer. The housing types used were the reference design original casing, un-housed, an indoor AP plastic, and an outdoor AP plastic. The housing variants are displayed in Figure 5.6 and the STA placements are shown in Figure 5.6a.

From the test results shown in Figure 5.7 it is apparent that there is very little to no variation in observed performance due to housing type. Even in Figure 5.6 e), which uses an outdoor AP plastic which is likely the most thick to protect against harsh environments and prevent degradation, there was no measurable difference in application layer performance. Since the closest client was at approximately 21 feet, it is possible that there are negative effects that the test did not capture. For example, if the distance to the clients were less, and there were enough clients at smaller spacings, it may be possible that the housing could add enough attenuation to reduce the range at which MCS boundaries appear. The antenna variation tests use finer granularity in client spacings, so a housing vs unhoused test is presented in Figures 5.20 and 5.21 which provides a slightly more detailed look at the unhoused vs indoor AP housing performance. Refer to that section for a a more detailed investigation of housing type effects.

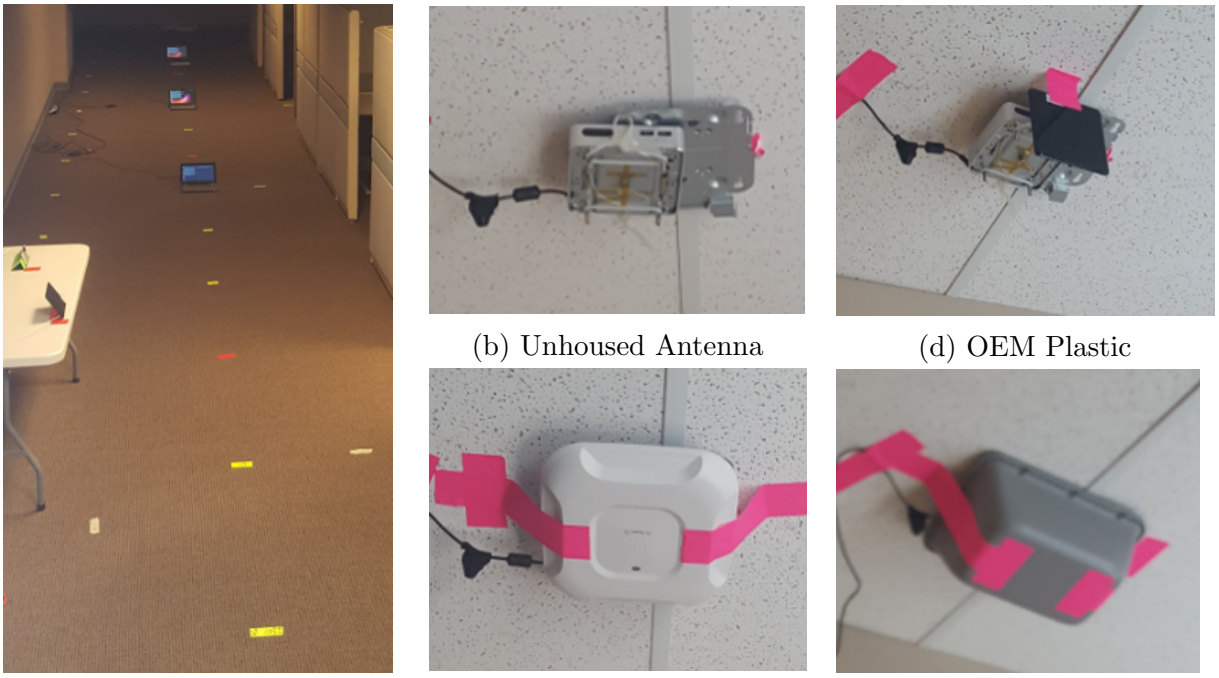

(a) STA Placement (b) Unhoused Antenna (c) Indoor AP Plastic (d) OEM Plastic (e) Outdoor AP Plastic

Figure 5.6: STA Orientation and Housing Plastic Types

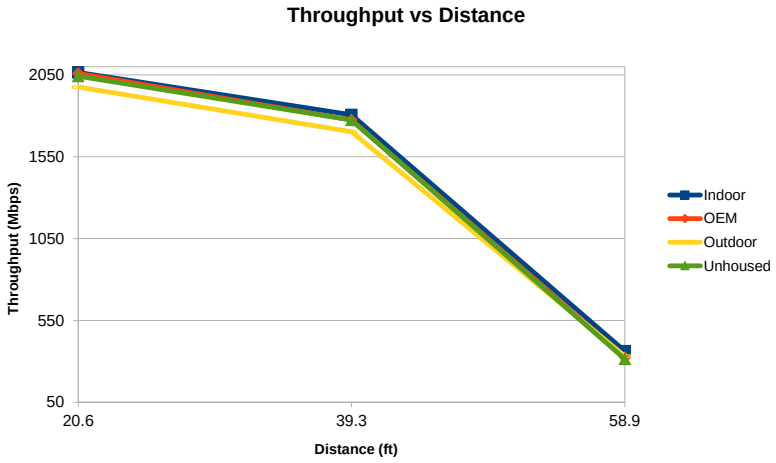

Figure 5.7: Downlink UDP Performance vs Housing Type

5.6 Station Orientation

The intent of the STA orientation testing was to determine the adverse effects STA orientation can have on link quality. With the exception of the later mobility tests, all tests have had the client oriented such that its maximum EIRP direction would be bore-sight with the AP. In the STA orientation testing, performance was profiled with a station orientation of -90, -45, 0, +45 and +90 degrees. Testing was conducted at distances of 11 ft and 25 ft to investigate the effects of distance. The average performance for each test is summarized below in Figures 5.9a and 5.9b.

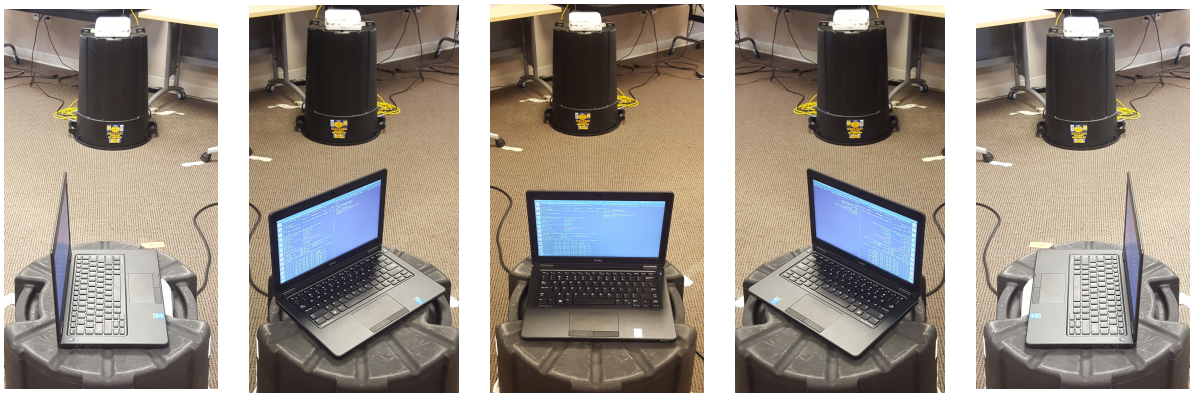

(a) -90° (b) -45° (c) 0° (d) +45° (e) +90°

Figure 5.8: Tested STA Orientations

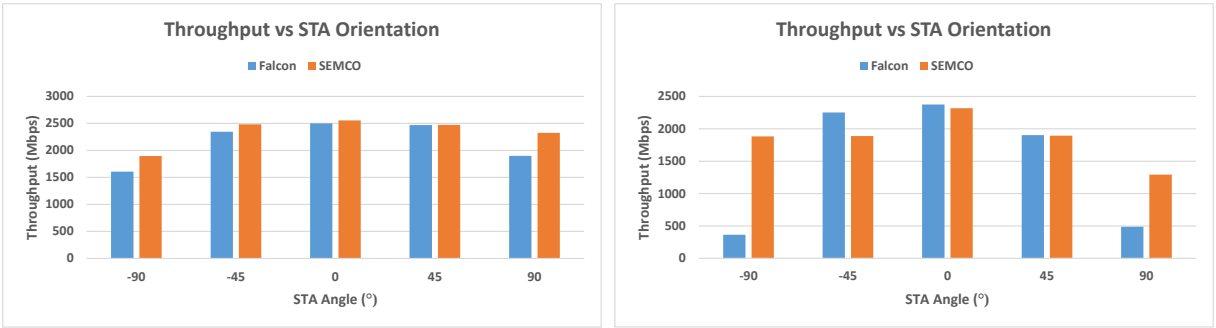

(a) Distance = 8 ft (b) Distance = 24 ft

Figure 5.9: Throughput vs STA Orientation vs Distance

From the orientation tests it is apparent that there is a benefit to having the antennas

aligned and sender and receiver. This is apparent in the chart as the highest throughput is measured at a zero degree angle (bore-sight), and performance degrades as the angle increases for both antenna styles. The degradation is less pronounced for the SEMCO (highly directional) AP antenna. It is likely that the use of additional transmit chains in this antenna allows the degradation to be partially mitigated by use of a high beamforming gain. Another interesting observation is that at the higher distance shown in Figure 5.9b the effects are more significant. This is likely an indication that the extra beamforming gain cannot mitigate the misalignment when the path length is much longer. For certain applications, like file-sync, or media download kiosks, antenna orientation and placement can be an important decision to favor the common case and leverage the highest performance offered by 802.11ad.

5.7 Obstruction

Another important consideration for wireless performance and reliability is obstruction testing. In realistic scenarios, it is very common to have objects blocking LOS between AP and client devices. Furthermore, for devices like smartphones which might be in the users pocket, it is important to consider the effects of human blockage. The goal of this suite of tests was to identify which materials 60 GHz signals could penetrate. The obstruction materials tested included a metal cabinet, a wooden door, a piece of ceiling tile, and a human. Many people are pessimistic about 802.11ad because they assume the link has to be completely unobstructed. These tests aimed to clarify what exactly the capabilities are.

Building Materials

The link could not be maintained through the metal cabinet as the metal acted as an antenna and absorbs most of the 60 GHz signal. The ceiling tile had no measurable affects on the application layer performance when introduced into the link. According to [46] ceiling tiles absorption coefficients add approximately 1.13 dB/cm of attenuation at 60 GHz. Since ceiling tiles are less than 2 cm thick, the attenuation was not noticeable enough to measure without link level instrumentation to accurately profile RSSI. Glass exhibits approximately 4 dB/cm of attenuation which is higher than the 1.3 dB/cm of attenuation added by a wooden door, however the thickness of single pane indoor glass is less than a centimeter. A 5 cm door thickness could add up to 6.5 dB of attenuation which is more significant than the ceiling tile or thin glass pane. The attenuation added by ceiling tiles and glass is lower than the increased FSPL due one foot increased endpoint separation.

This slightly reduces the spatial boundaries of certain MCS rates but is not significant in practice. Drywall adds attenuation of 0.09 dB/cm and glossed drywall adds 0.6 dB/cm which is essentially negligible. Other building materials found on outer walls and floors add significant attenuation. Brick wall adds attenuation of 14.7 dB/cm, cement cinder block adds 11.3 dB/cm, and cement backer board is the highest at 17 dB/cm, likely due to dual plexiglass coating (8.6 dB/cm)[46]. The attenuation of these materials prevents WiGig service from being maintained through thick walls, floors and building boundaries but also allows higher frequency re-use and lower chance of eavesdropping.

Wooden Door

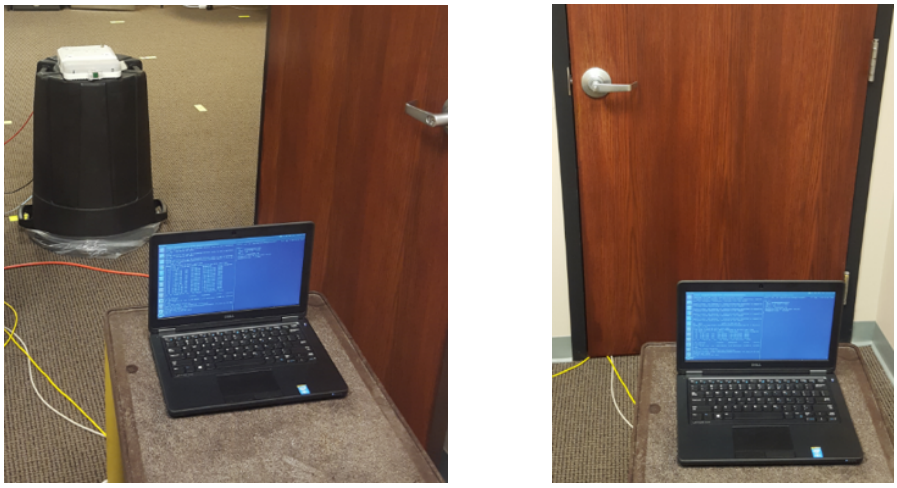

(a) Control Baseline (Door Open) (b) Obstructed (Door Closed)

Figure 5.10: Wooden Door Test

From the results presented in Figure 5.11, it is clear that strong beamforming gain can allow a WiGig link to be maintained through a wooden door. For the unobstructed baseline, the client and AP are close enough that MCS rates 10 and 12 can be obtained providing PHY rates up to 4 Gbps. Once the door is closed, the PHY rate drops to between 2.2 and 2.5 Gbps at MCS 8 and 10. While this is a fairly significant drop, maintaining over two Gbps of throughput through a door with no other propagation paths is significant, and indicates that 60 GHz Wi-Fi does not simply obey quasi-optical propagation behaviors. The extra attenuation caused by the door does cause the SNR to decrease, but not enough to cause a link failure, and the throughput is still significantly higher than even the fastest 802.11ac configuration.

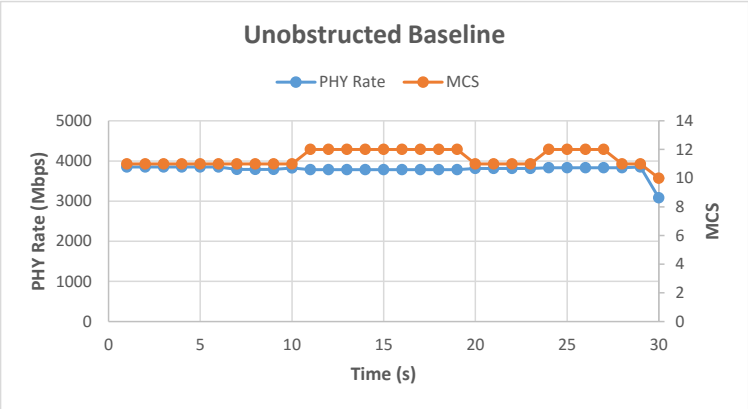

(a) Control Test (Door Open)

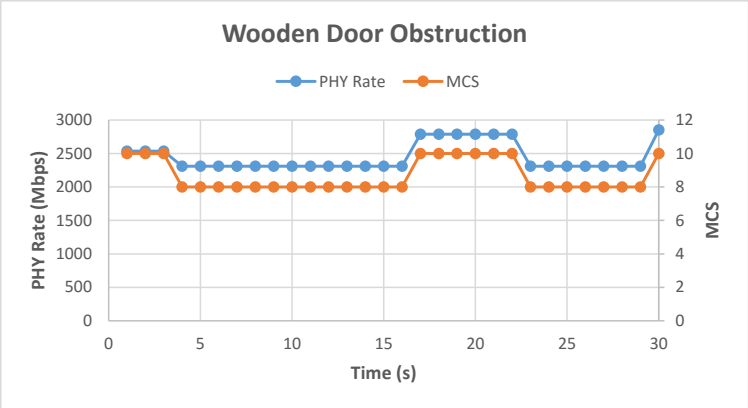

(b) Obstructed Test (Door Closed)

Figure 5.11: Door Characterization

Human Obstruction

In the case of human obstruction, two types of blockage were of interest: persistent and transitory. Both tests began unobstructed (10 seconds) but shortly thereafter the link was obstructed by a human obstacle. In the former, the human would remain for the duration of the test (20 seconds). In the latter tests, the obstruction was transitory to investigate if once the blockage was relieved the BFT could quickly re-obtain higher data rates. Another variable of interest was how the environment affected the links ability to avoid obstacles. For example, in a small closed in room the link can be maintained by leveraging one of the

many multiple propagation paths between sender and receiver (e.g. a first order reflection off a nearby wall). In a large open room, the walls are not necessarily close enough for these multipaths to be used in a constructive manner. Thus persistent and intermittent human obstruction tests were ran in both the conference room shown in Figure 4.9, and the open office shown in Figure 4.10.

Meeting Room

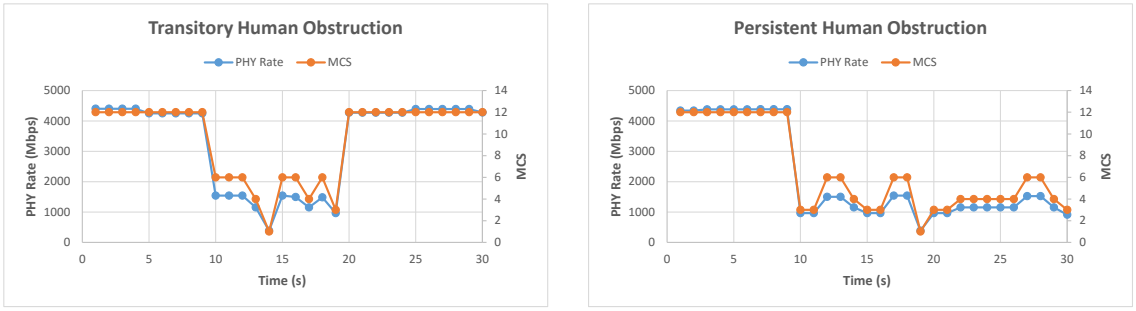

(a) Transitory (b) Persistent

Figure 5.12: Human Obstruction (Meeting Room)

Open Office

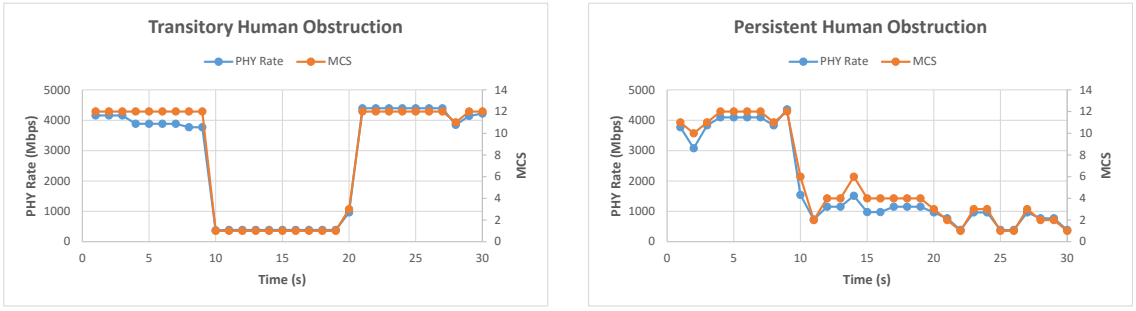

(a) Transitory (b) Persistent

Figure 5.13: Human Obstruction (Open Office)

The results are as to be expected. When an obstruction blocks LOS between endpoints several successive failed transmissions cause the link to degrade to a more robust and

reliable modulation scheme. This is apparent in the quick drop in performance at t=10s in Figures 5.12 and 5.13. In the transitory obstruction tests, once the obstacle is removed, the link quickly re-trains to its original optimized rate. The time between low-rate and high-rate is less than a second, indicating that the link can optimize and adjust itself extremely quickly. Another very interesting observation is the difference in fall-back rates in the meeting room compared to the open office. As expected, in the small meeting room, when the link is obstructed the signal is reflected off a nearby wall to the client, allowing usage of a first-order reflection off with only a minor increase in path length. In the open office space, the nearest wall is approximately 30 ft out. In order to obtain a suitable reflection, the path length increases substantially resulting in increased fading and absorption which causes the link to downgrade to a much lower data rate. Multi-path is typically a barrier for wireless communication but in the presence of transitory obstructions, these tests have shown that multiple propagation paths can be beneficial for link reliability once obstructions are introduced. The obstructed throughput in the meeting room is still over 1 Gbps, whereas in the open office it must degrade to nearly the lowest data rate, around 200 Mbps. In summary, it appears that 802.11ad links can be maintained even in the presence of persistent obstructions. The rate shifting algorithms used quickly converge to appropriate data rates to maintain communication. In large open office spaces, it is apparently more difficult for the link to mitigate obstruction degradation as there are less surfaces to reflect the signal off of. Even then, the 200 Mbps seen in the open office space is very high performance and is suitable for many applications.

5.8 Mobility

In all of the tests conducted up until now both the client and AP have been stationary. With the exception of minor variation in O2 contcentration, humidity, and multipath due to human motion, the channel is sufficiently static. However, in many realistic scenarios the client may be moving around which creates additional difficulties for wireless, especially due to the high directionality of 802.11ad links. In the mobility tests points on a straight line where selected and static-link performance was measured. The client was then moved between the points at various speeds to evaluate the effects of mobility on link quality. For these tests the client was placed on a moving cart, and the AP was mounted on the ceiling exactly in the middle (at 28 feet). The test was run for five seconds with the client static, then the cart was pushed from end to end at an approximately uniform speed. The end time of the test was recorded after each run. The total distance traveled was always 56 feet which allowed post-computation of the cart speed. The data was processed to remove the

static portions of the test. For each speed (low, medium, high) three tests were run from left to right, and from right to left, to eliminate biases due to the five second link training interval at the start. The speeds were approximately a walking pace, a quick walk, and a fast jog. Throughput was plotted as a function of distance including the static baseline tests. The results are presented in Figures 5.14a, 5.14b, and 5.14c. The average velocity, performance and standard deviation are summarized in Table 5.4.

Table 5.4: Performance vs Client STA Speed

$\mathbf{V_{ave}}$ **(ft/s)**	$\mathbf{TP_{ave}}$ **(Mbps)**	$\mathbf{StdDev_{ave}}$ **(Mbps)**
0 (static)	1686.4	496.89
1.5545	1597.8	551.42
2.6599	1566.4	549.40
4.3316	1381.7	609.42

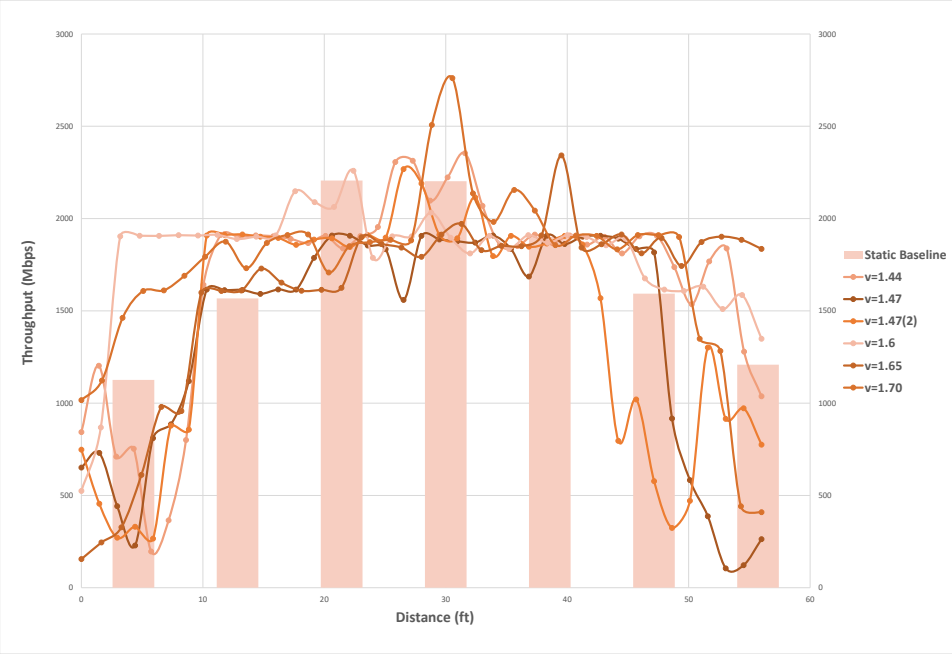

(a) Low Speed Test

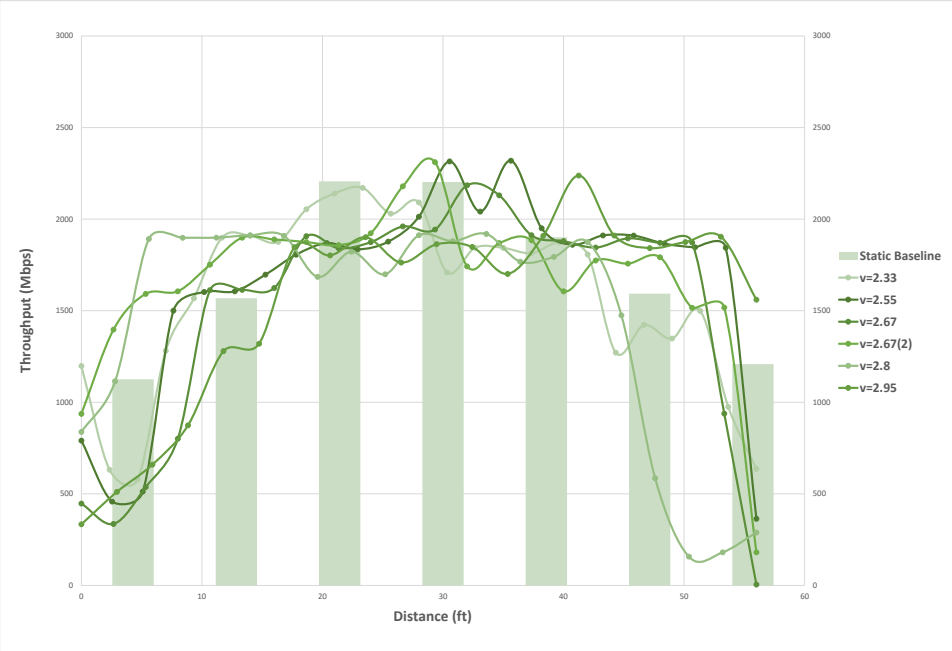

(b) Medium Speed Test

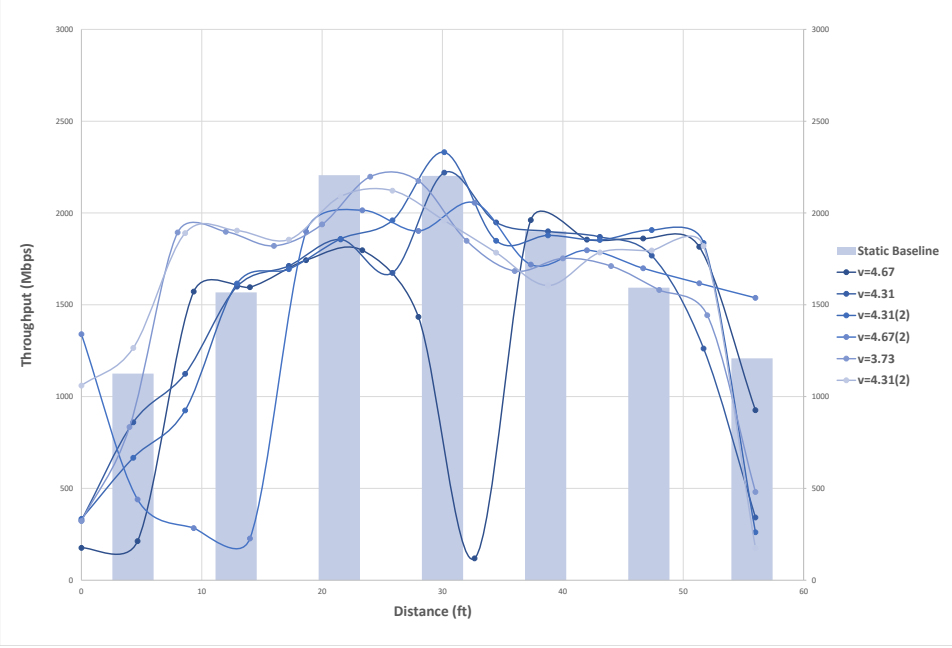

(c) High Speed Test

Figure 5.14: Performance vs Speed

From the average data in Table 5.4, it is apparent that there is a steady decrease in performance and reliability as station velocity increases. This is to be expected as there is slightly more overhead due to usage of extra beamforming training intervals to mitigate the effects of the dynamic link. As expected, since the AP is closest to the client when d=28 ft, the link performance increases until that point and decreases afterwards. In the lower speed tests the highest throughput peaks are more pronounced at d=28 ft compared to the higher speed tests. This is likely because as the client speed increased it was more difficult for the link to obtain higher data rates such as MCS 11 and 12. Despite the fairly high speed in the fastest test, the link can maintain application level throughput of over 1.5 Gbps between distances of 10 and 46 feet. This result indicates that in the presence of only mobility high LOS link performance can still be maintained with 802.11ad, despite a heavy reliance on directional transmission. Qualitatively the beamforming training converges quickly enough such that relative motion of handheld devices is not a major concern for link performance. The highest throughput applications of 802.11ad, such as wireless backhaul, uncompressed video streams, sync-and-go, would typically be done between static endpoints (e.g. above ceiling tiles, between building rooftops, or between a media centre and a smart TV). In these situations the link would be sufficiently static that mobility degradation would not be a concern. In the case of portable devices such as smartphones it is unlikely that the station would be moving as quickly as in the medium and high speed tests. The empirical throughput observed in the presence of client mobility around 2 ft/s indicates that 802.11ad would still be highly beneficial for small-cell applications even if the client were not stationary.

5.9 Long Hall Rate-vs-Range and Bridge Testing

The primary intent of the long hall LOS RvR tests were to determine the spatial boundaries where specific MCS rates could be achieved. Another area of interest was the maximum range at which a reliable link could be achieved. Detailed measurements were made using both the highly directional AP-style antenna, and lower power client-style antenna to see if there was a practical benefit of using the directional antenna for bore-sight performance. Long hall RvR was collected for both an AP to STA configuration but also an AP to AP bridge. Configuring the PoC platform as both ends of the link allows measurement of send and receive bottlenecks and asymmetry in uplink/downlink performance.

The long hall bridge results presented in Figure 4.12 indicate several interesting behaviors. When comparing the AP and client antennas it is apparent that MCS boundaries differ (shown in blue). It appears that either antenna could maintain MCS 11 and 12 up

to 16 feet, however from 16 to 40 feet the AP-style antenna could achieve MCS 10 whereas the client-style antenna immediately dropped to MCS 8. The use of extra transmit chains on the AP-style antenna allows a notable increase in the range of some of the higher order MCS (10 and 11). At distances below 16 ft, and above 40 ft there was no noticeable difference in achievable MCS and TP.

The next notable difference is asymmetry in application layer TP between uplink and downlink tests at close range. As mentioned, the PoC platform suffers by some processing limitations capping its goodput to 2 Gbps. On the uplink, this was mitigated by adding a wired endpoint to terminate some of the traffic. The offloading of some of the data processing allowed a 20% increase in performance (2 Gbps vs 2.4 Gbps). By comparing the goodput (grey curve) between Figures 5.16a, 5.16c and Figures 5.16b, 5.16d it is apparent at the higher MCS rates (10, 11, 12), the Layer-4 performance for the uplink was higher. Due to the extension of the MCS boundaries of the AP-style antenna, the asymmetry is apparent until 40 ft, whereas for the client-style antenna the offloading only benefits up to 20 ft.

The last important observations are made between the long hall client RvR, and the long hall bridge RvR. For the most part, the AP to STA link performed the same as the AP bridge. The client laptops features high performance Intel i7s which removed the uplink/downlink asymmetry since the downlink testing did not suffer from the aforementioned bottleneck. One difference was that the client laptop could obtain MCS 10 on uplink from distances between 24 ft and 40 ft. This could be due to random differences in link spacing (experimental error), or minor variation in station orientation. Lastly, in three of the tests the furthest point (d=160 ft) dropped off below MCS 8. This could be due to the fact that the AP bridge testing used unhoused antennas at both ends. In the client long hall tests, the client antenna modules were beneath the laptop casing. The very minor attenuation added by the plastic casing becomes apparent at MCS boundaries with fine granularity station spacing, or at the far end of the service area.

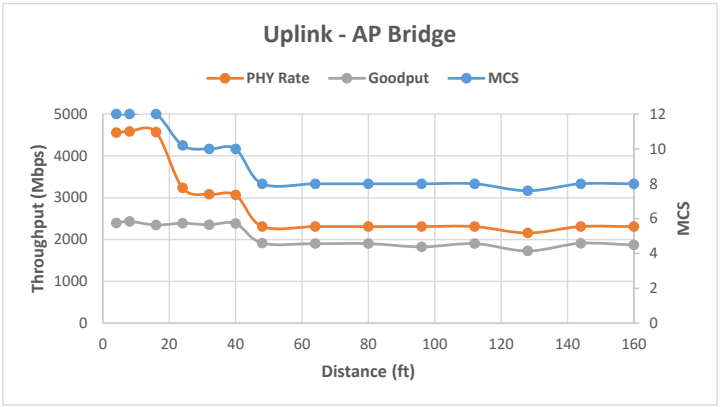

(a) AP Antenna - Uplink

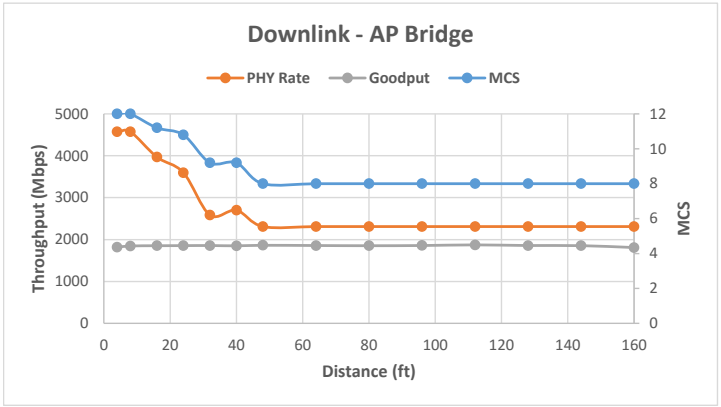

(b) AP Antenna - Downlink

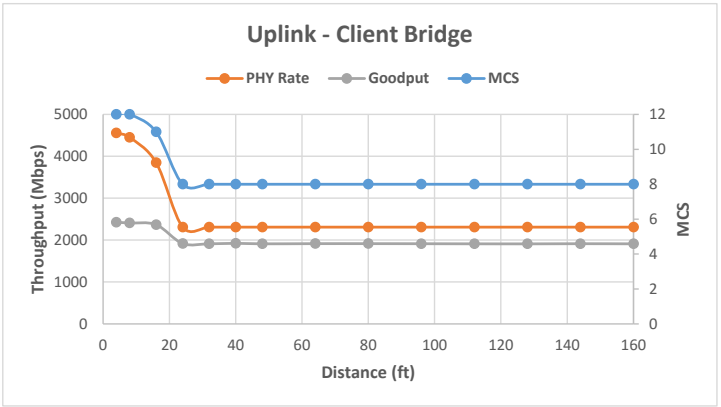

(c) Client Antenna - Uplink

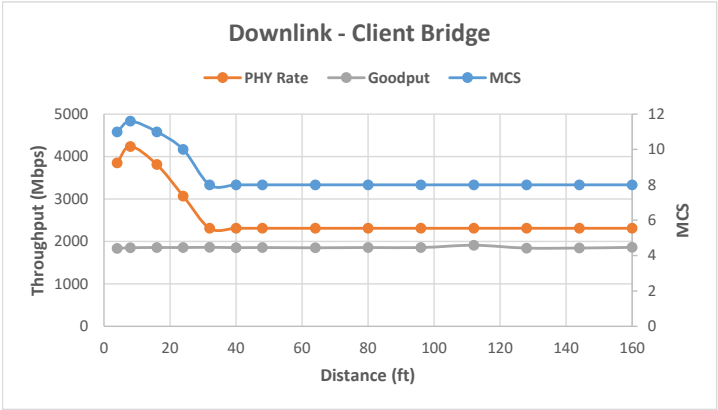

(d) Client Antenna - Downlink

Figure 5.15: Long Hall Bridge Tests

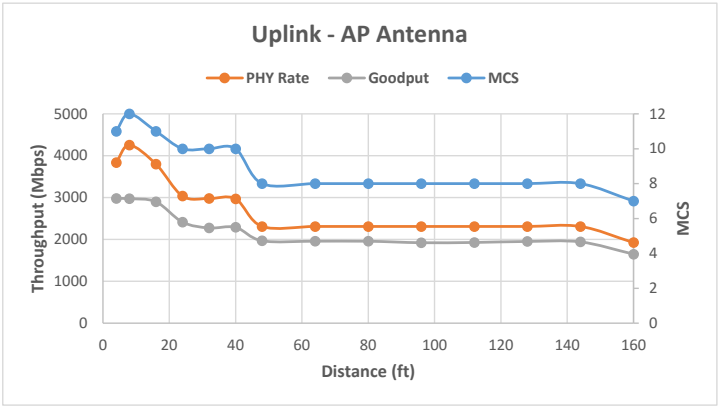

(a) AP Antenna - Uplink

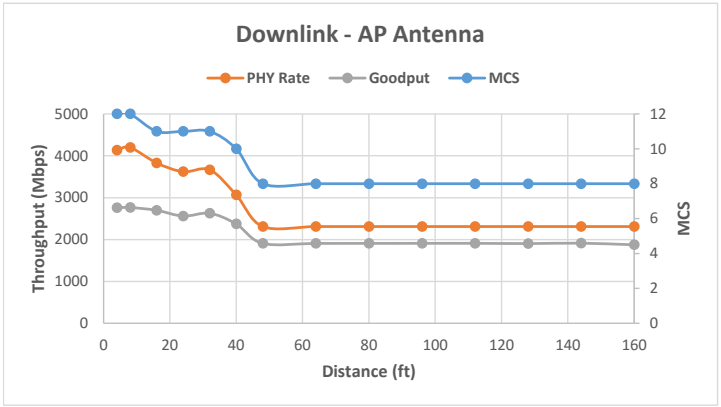

(b) AP Antenna - Downlink

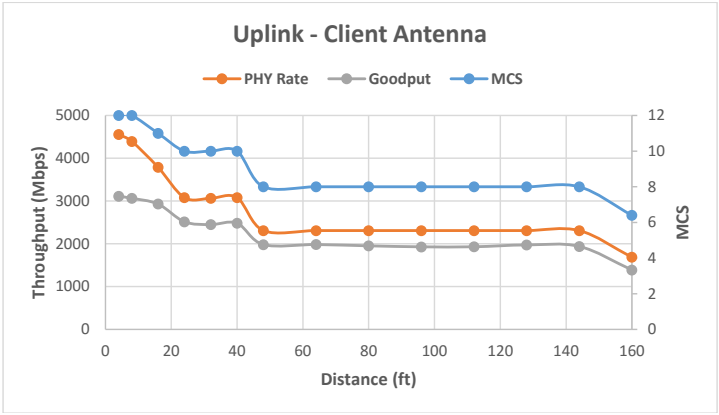

(c) Client Antenna - Uplink

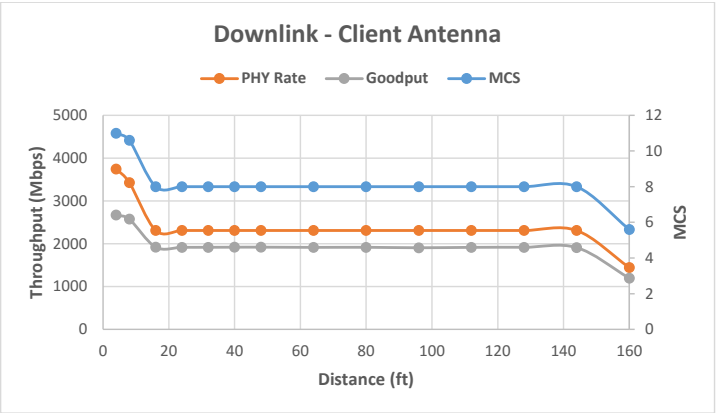

(d) Client Antenna - Downlink

Figure 5.16: Ubuntu STA (Client Antenna) Long Hall Tests

5.10 Open Office Antenna Characterization

The intent of this suite of tests was to profile the practical design considerations when choosing an AP industrial design. Specifically, these tests were aimed at identifying if the AP orientation had significant effects on the service area for 60 GHz WiFi. Since the extension module is side-car mounted, as shown in Figure 4.8g, profiling the service area could potentially identify 'shadowing' effects on the opposite side of the module and other regions of nodal and anti-nodal interference. Another design consideration of interest was whether the antenna orientation has a significant effect on performance and range, e.g. whether the antennas maximum EIRP direction should be orthogonal or parallel with the ground plane. It is also valuable to see the difference in coverage patterns between various antenna topologies (e.g. highly directional vs quasi omni directional antenna styles). This is important to investigate for the very highly directional antenna, if the performance drastically diminishes at high angles to the antennas bore-sight. Finally, the 'periscoped' antenna module shown in Figure 4.8c was tested to see if having an alternative antenna placement could alleviate any potential shadowing effects.

AP Orientation - Shadowing

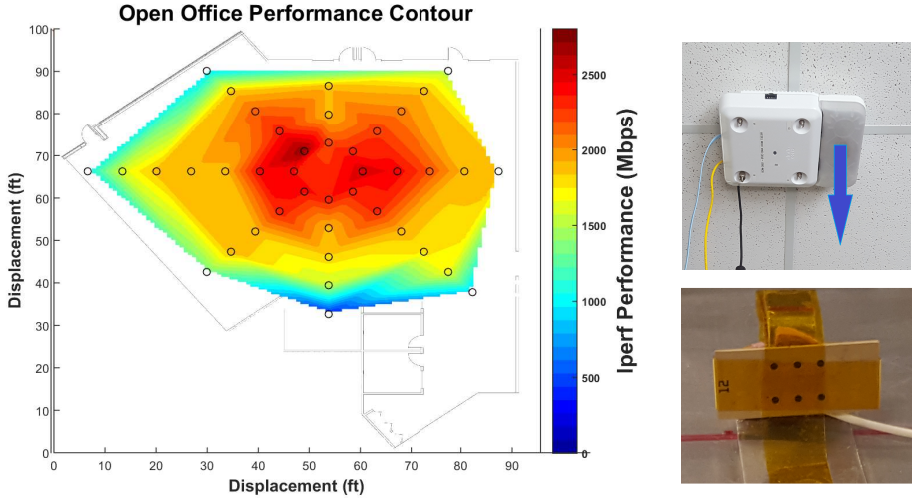

(a) Topographical Throughput Map (b) AP and Antenna Configuration

Figure 5.17: AP Oriented Right

From Figures 5.17 and 5.18 it is apparent that when the 60 GHz RF front-end is side mounted the multi-Gbps service area (dark red) is slightly extended on the module side.

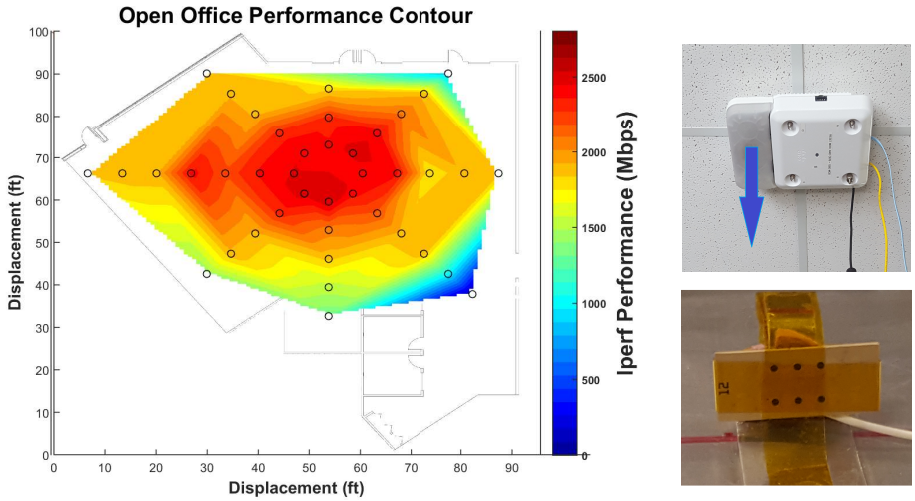

(a) Topographical Throughput Map (b) AP and Antenna Configuration

Figure 5.18: AP Oriented Left

In Figure 5.17 the dark red extends three test locations to the right of the AP (approx. 25 ft), and only two to the left (approx 18 ft). When the module is on the left side, as in Figure 5.18, the dark red extends four test locations to the left (approx. 33 ft), and roughly 25 ft to the right (with a small anti-nodal region appearing on the shadowed side). This suggests that the AP should be oriented with the 60 GHz front-end as close as possible to the majority of clients or those with highest demands. Furthermore, since the dark red service area does not extend to either of the open office boundaries (approx. 100 ft), with this particular client-style antenna it is optimal to place the AP in the center of the room. It is likely that if the AP at either of the far sides of the room, that there would be more pronounced degradation on the opposite side of the room. The periscoped antenna shown in Figure 5.19 alleviates all of the shadowing caused by the placement of the EM. By allowing the antenna array to protrude downwards from the AP the shadowed region that appeared on the opposite sides of Figures 5.17 and 5.18 are removed. Furthermore, nearly all test points exhibited much higher average performance than any other antenna design. It is important to note that the periscoped module was not housed, so a small amount of the increased performance could be due to slightly lower housing attenuation. For future 60 GHz front end design these results suggest that the antenna should be displaced slightly from the AP housing so shadowing does not occur at extreme angles.

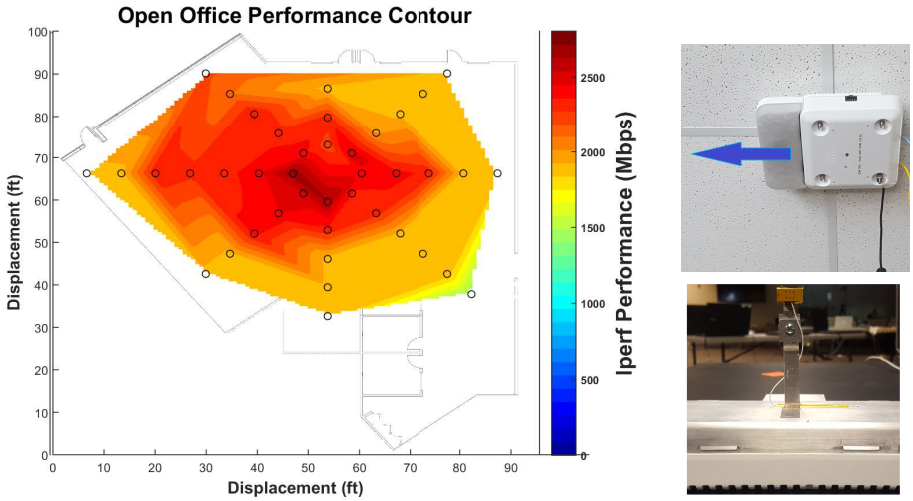

(a) Topographical Throughput Map (b) AP and Antenna Configuration

Figure 5.19: AP Oriented Left - Periscoped Antenna - Perpendicular with Ground Plane

Housing Attenuation

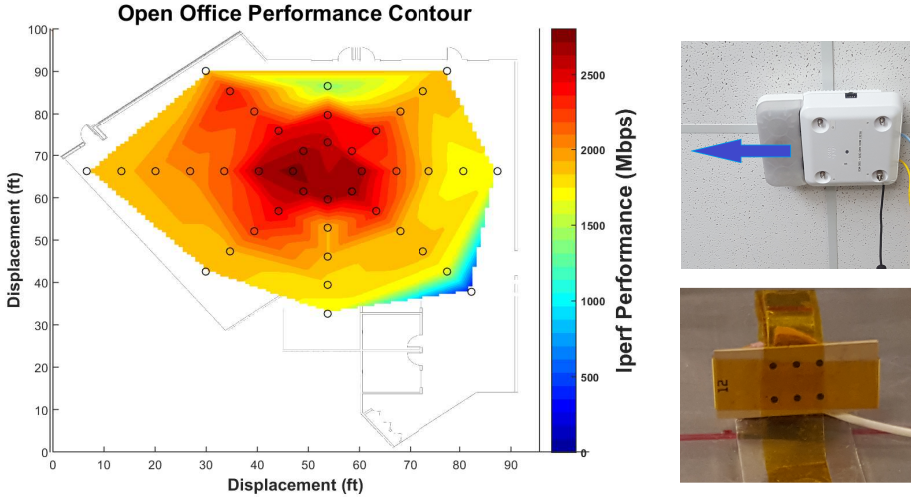

(a) Topographical Throughput Map (b) AP and Antenna Configuration

Figure 5.20: Housed Falcon Antenna

The housing test results shown in Figures 5.20 and 5.21 show that while many test locations are unaffected specific edge locations may be achieve a minor performance increase

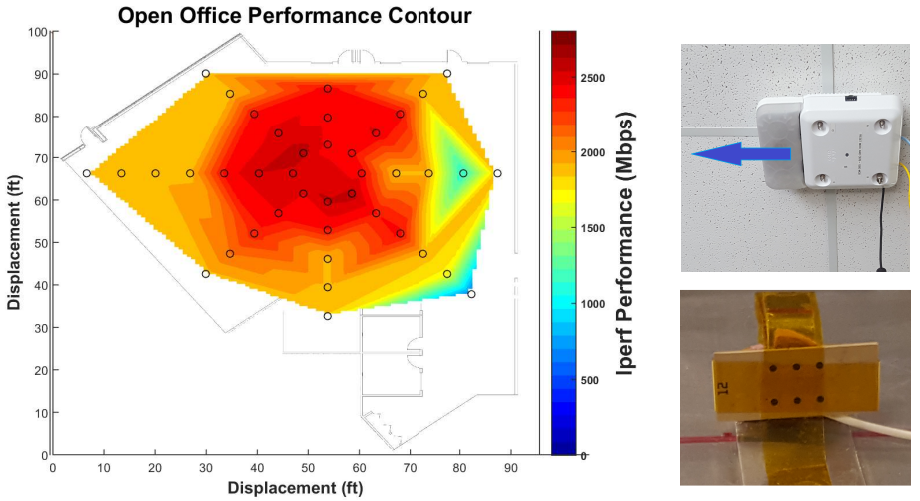

(a) Topographical Throughput Map (b) AP and Antenna Configuration

Figure 5.21: Unhoused Falcon Antenna

by removing antenna housing. The points along a vertical line intersecting the center of the diagram typically achieve lower performance, because they are at a 90°angle with the antennas bore-sight. Conversely, the points along a horizontal line intersecting the AP bore-sight (i.e. 0°and 180°) exhibit improved range and performance. This is also indicated by the oval shaped service area suggesting that the rectangular patch antennas used in these designs allow a less omni-directional service area in comparison to traditional dipole antennas. Comparing Figures 5.21 and 5.19 suggests that the use of the periscope in fact did remove the shadowed region on the far right as both antenna styles were unhoused and oriented the same.

Directional Antenna Coverage

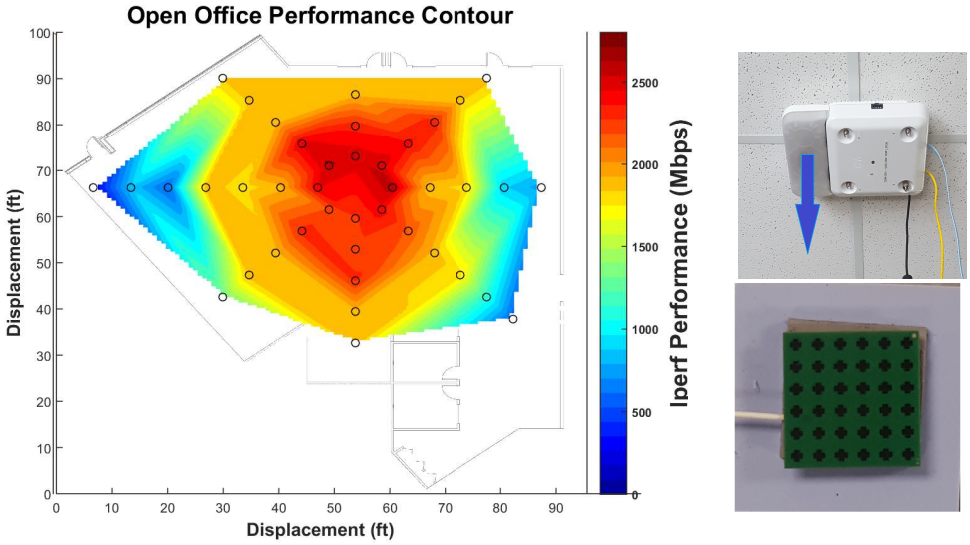

(a) Topographical Throughput Map (b) AP and Antenna Configuration

Figure 5.22: Directional Antenna - Flush with Ground Plane

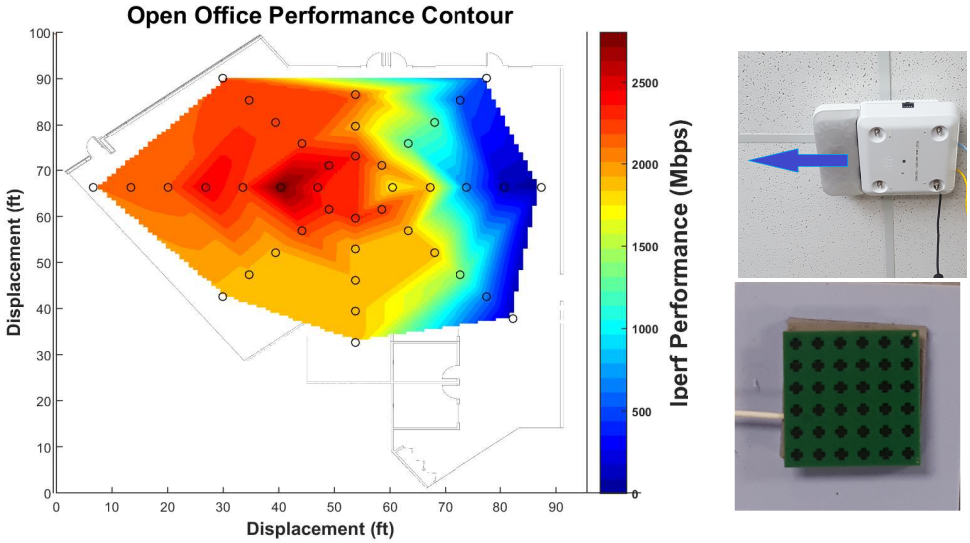

(a) Topographical Throughput Map (b) AP and Antenna Configuration

Figure 5.23: Directional Antenna - Perpendicular with Ground Plane

The AP-style antennas service area shown in Figures 5.22 and 5.23 exhibit more significant anti-nodal regions than the client-style antenna. The added antenna gain due to the use of extra transmit chains does not have pronounced effects when all test locations are nearby. In the left and right edges of Figure 5.22 performance drops off slightly below 1 Gbps, which is lower than observed with the client-antenna. This is due to the fact that the directional antenna was oriented flush with the ground plane. This causes its maximum EIRP direction to be directed at the ground 10 ft below. This diminishes the benefit of the higher power antenna as the use of a single polarization prevents it from providing high performance at higher angles. When the maximum EIRP direction is perpendicular with the ground plane as in Figure 5.23, the anti-nodal regions are removed on the side where the antenna is directed. However, on the shadowed side, the cut-off becomes even more pronounced. Since the AP-style antenna is intended to be used in a diversity array, it is clear that even two of these antennas could provide very high service for the open office space (one directed left, one directed right). This would allow the benefit of the extended range and mitigate the reduced service area allowing a broader high-performance coverage area. It would be interesting to test this configuration again, but with a periscoped AP-style antenna to determine if the performance could be improved at high angles (e.g. 180°).

5.11 Spatial Re-Use

Another concern with wireless technologies is multi AP interference. Multi AP interference occurs when multiple service areas from adjacent networks overlap. When the two service areas are on the same channel STAs may sense neighboring transmissions during CCA and back off the medium. Contention for the channel could be extremely high reducing cumulative throughput. When the two service areas are on adjacent channels, any overlap in signal bandwidth can be picked up as interference by receivers resulting in a lower SNR causing lower transmission data rates to be used. The intent of these tests were to identify the extent of multi AP interference when they are placed on the same channel (co-channel) or adjacent channels. The study presented in [33] proposes a spatial-reuse factor which is defined as:

$$\beta = \frac{A + B}{C} \tag{5.1}$$

where:

A = performance of Link 1 when both links are active
B = the performance of Link 2 when both links are active
C = the average performance of Link 1 and 2 and run independently

In the ideal case, the spatial reuse factor β works is exactly two, i.e. the links can be run simultaneously without any reduction in cumulative performance compared to independent operation. The intent of the spatial re-use testing was to establish estimates for β at various link spacings, when the links overlap, and when they both operate on the same channel (co-channel), or separate channels (adjacent channel). This is one way to characterize the efficacy of the spectrum mask shown in Figure 2.3. In co-channel tests, β indicates the performance reduction effects of the MAC layer EDCA algorithms, and provides insight into how the use of highly directional beams reduces signal interference and unwanted eavesdropping.

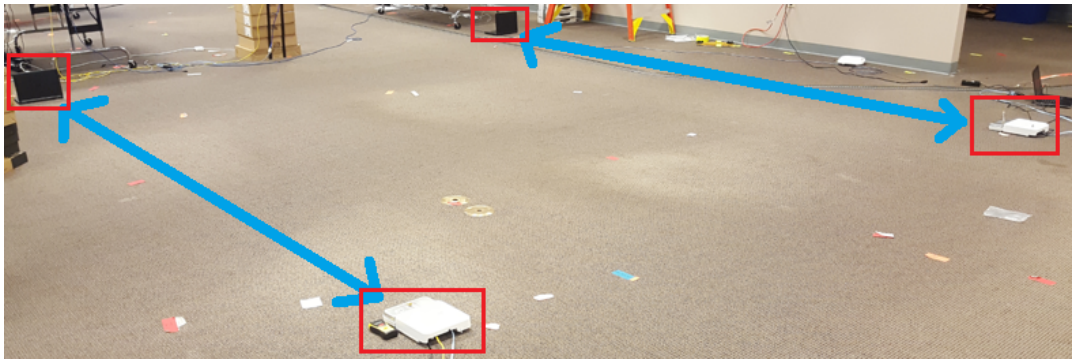

Figure 5.24: 16 ft Link Separation

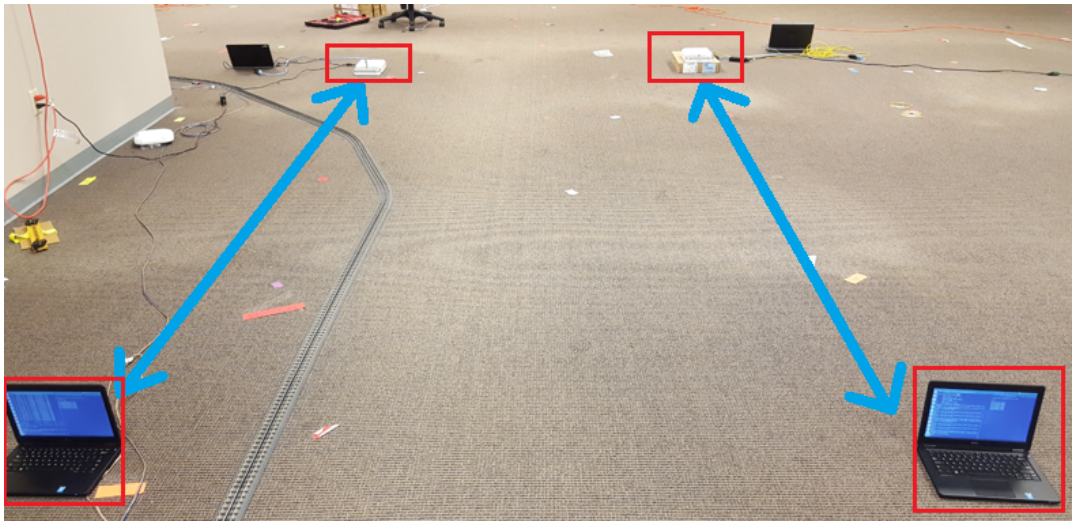

Figure 5.25: 8 ft Link Separation

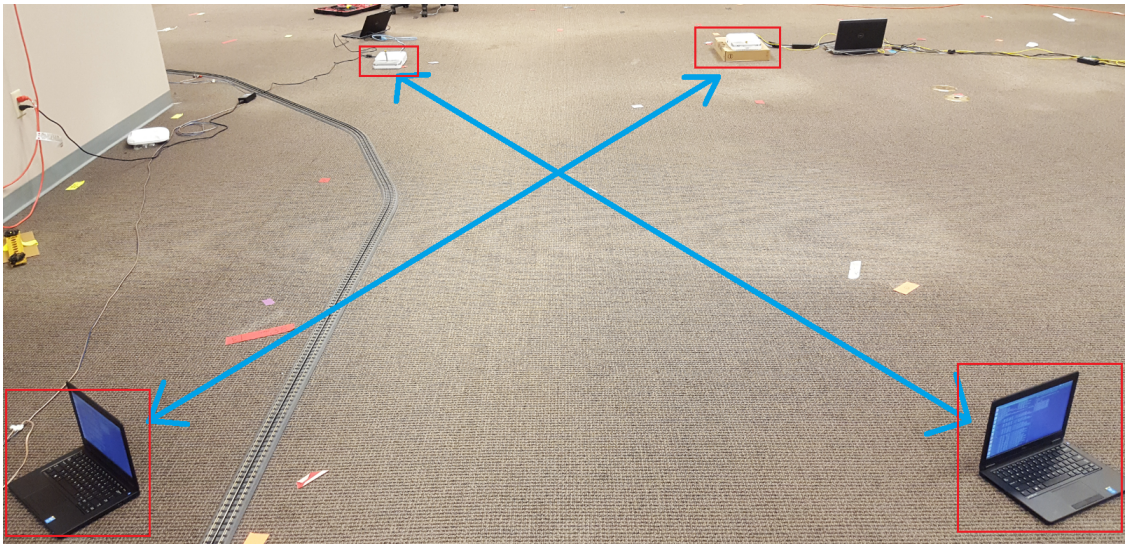

Figure 5.26: Cross Link Setup

5.11.1 Adjacent Channel Interference

The results shown in Tables 5.5–5.10 (Note - TP/loss in Mbps) indicate that with these test platforms there was no ACI between Channel 1 (f_c = 58.32 GHz) and Channel 2 (f_c = 60.48 GHz) due to the spectral mask defined in the standard. There was no measurable degradation due to adjacent channel network overlap even in the cross-link tests.

Traditional low and mid end Wi-Fi repeaters and extenders use the same channel as the network they are repeating. This essentially halves the wireless capacity since they must first receive the packet and then repeat it (doubling the total number of transmissions). High-end repeaters can use non-overlapping channels which allows the capacity to be unaffected with a small increase in latency. Due to the high 60 GHz channel availability in most regions it is feasible to have very high network density and co-existence using any combination of the four non-overlapping channels. This greatly improves the spectral efficiency and can enable usage of wireless repeaters to extend the range of WiGig, and to provide LOS service around blockages by using intelligently placed relays.

Table 5.5: ACI (d=16 feet) - Cisco EM

	Test Case	Link1	Link2			
Solo	Uplink	2360	2470			
	Downlink	1567	1554	Loss		β
Multi	Up-Up	2469	2282	108	-188	1.97
	Dn-Dn	1579	1451	12	-103	1.94
	Up-Dn	2483	1546	123	-8	2.06
	Dn-Up	1607	2333	40	-137	1.95

Table 5.6: ACI (d=16 feet) - Reference Design

	Test Case	Link1	Link2			
Solo	Uplink	1766	1949			
	Downlink	1626	1546	Loss		β
Multi	Up-Up	1831	2020	65	71	2.07
	Dn-Dn	1683	1505	57	-41	2.01
	Up-Dn	1798	1569	32	23	2.03
	Dn-Up	1600	1921	-26	-28	1.97

Table 5.7: ACI (d=8 feet) - Cisco EM

	Test Case	Link1	Link2			
Solo	Uplink	2280	2405			
	Downlink	1586	1497	Loss		β
Multi	Up-Up	2502	2541	223	137	2.15
	Dn-Dn	1596	1570	10	73	2.05
	Up-Dn	2421	1476	141	-21	2.06
	Dn-Up	1543	2412	-43	8	1.98

Table 5.8: ACI (d=8 feet) - Reference Design

	Test Case	Link1	Link2			
Solo	Uplink	1768	1905			
	Downlink	1461	1506	Loss		β
Multi	Up-Up	1778	1858	10	-47	1.98
	Dn-Dn	1504	1501	43	-5	2.03
	Up-Dn	1756	1505	-12	-1	1.99
	Dn-Up	1566	1902	105	-3	2.06

Table 5.9: ACI (Cross Link) - Cisco EM

	Test Case	Link1	Link2			
Solo	Uplink	2464	2274			
	Downlink	1578	1571	Loss		β
Multi	Up-Up	2447	2178	-17	-96	1.95
	Dn-Dn	1590	1358	12	-213	1.87
	Up-Dn	2042	1528	-422	-43	1.77
	Dn-Up	1565	2177	-13	-97	1.94

Table 5.10: ACI (Cross Link) - Reference Design

	Test Case	Link1	Link2			
Solo	Uplink	1711	1725			
	Downlink	1723	1561	Loss		β
Multi	Up-Up	1709	1708	-2	-17	1.99
	Dn-Dn	1722	1524	-1	-37	1.98
	Up-Dn	1702	1532	-9	-29	1.98
	Dn-Up	1671	1707	-52	-18	1.96

5.11.2 Co-Channel Interference

The results shown in Tables 5.11–5.16 (Note - TP/loss in Mbps) indicate that with these test platforms there was measurable degradation due to CCI. Comparing Table 5.11 to Table 5.16 it is apparent that the CCI increased as link spacing decreased. At spacings of 16 feet and 8 feet the spatial re-use factor was over one suggesting that the total capacity still increased when adding a second link. This is likely due to baseband processing limitations on either link resulting in channel usage being lower than 100% with only a single link active. When a single link gets saturated the nodes would backoff the medium resulting in a silent period which could allow a TXOP for the second link.

Another noticeable affect was that during the cross link tests sometimes the clients would disconnect from the BSS. This was most likely due to the AP not being able to send its beacon frame due to CCA detecting the medium as busy. This made it increasingly difficult to run tests without one or both clients disconnecting due to lack of beacons. Another noticeable difference was that the spatial re-use was measurably lower for the reference designs compared to the PoC. One potential reason for this is due to hardware revisions and driver versions. The PoC was engineered a full year after the reference designs and featured the latest wil6210 driver and firmware. It is likely that the MAC implementation had been slightly refined, or that the hardware was improved resulting

in less spurious transmission, narrower beam-width, or other advancements which would result in improved performance.

Despite there being measurable CCI in this test suite these link spacings and load demands were unrealistic. The networks are almost entirely overlapping. It is uncommon for that to be a necessity in a real world deployment. Typically they can be on the opposite ends of rooms where the high directionality and narrow beam-width at 60 GHz allows lower adjacent and co-channel interference. Furthermore during these tests both links attempted to obtain the maximum possible throughput, which is also atypical. Even high resolution video streaming to multiple clients would not consume nearly 2 Gbps of bandwidth. Due to both links attempting to send at very high rates collisions occur frequently and contention is higher resulting in lower total capacity. In summary, there is some measurable CCI at very high network densities and high load demands, but realistic deployments would not have such stringent bandwidth and spacing requirements.

Table 5.11: CCI (d=16 feet) - Cisco EM

	Test Case	Link1	Link2			
Solo	Uplink	2072	2018			
	Downlink	1534	1619	Loss		β
Multi	Up-Up	1199	1451	-872	-567	1.30
	Dn-Dn	943	843	-592	-776	1.13
	Up-Dn	1691	656	-380	-963	1.27
	Dn-Up	835	1669	-699	-350	1.41

Table 5.12: CCI (d=16 feet) - Reference Design

	Test Case	Link1	Link2			
Solo	Uplink	1700	1765			
	Downlink	1692	1538	Loss		β
Multi	Up-Up	1192	1371	-508	-394	1.48
	Dn-Dn	1070	808	-622	-730	1.16
	Up-Dn	1103	907	-597	-631	1.24
	Dn-Up	1225	1626	-467	-139	1.65

Table 5.13: CCI (d=8 feet) - Cisco EM

	Test Case	Link1	Link2			
Solo	Uplink	1985	1920			
	Downlink	1579	1575	Loss		β
Multi	Up-Up	1209	1017	-776	-902	1.14
	Dn-Dn	906	920	-673	-655	1.16
	Up-Dn	1308	783	-676	-791	1.18
	Dn-Up	796	1195	-783	-725	1.14

Table 5.14: CCI (d=8 feet) - Reference Design

	Test Case	Link1	Link2			
Solo	Uplink	1748	1940			
	Downlink	1560	1450	Loss		β
Multi	Up-Up	772	936	-976	-1004	0.93
	Dn-Dn	911	705	-649	-745	1.07
	Up-Dn	1081	581	-667	-869	1.04
	Dn-Up	666	1057	-894	-883	0.98

Table 5.15: CCI (Cross Link) - Cisco EM

	Test Case	Link1	Link2			
Solo	Uplink	2064	1909			
	Downlink	1521	1554	Loss		β
Multi	Up-Up	1237	1159	-826	-750	1.21
	Dn-Dn	849	931	-673	-623	1.16
	Up-Dn	1073	905	-990	-649	1.09
	Dn-Up	883	1123	-638	-786	1.17

Table 5.16: CCI (Cross Link) - Reference Design

	Test Case	Link1	Link2			
Solo	Uplink	1649	1746			
	Downlink	1635	1545	Loss		β
Multi	Up-Up	849	730	-800	-1016	0.93
	Dn-Dn	874	481	-761	-1064	0.85
	Up-Dn	863	473	-786	-1072	0.84
	Dn-Up	713	743	-922	-1003	0.86

5.12 Multiclient Spectral Efficiency

Another concern with regards to wireless performance is spectral efficiency as the number of client stations increases. If there are many endpoints trying to use the medium the overhead caused by channel backoff and re-transmission can result in an overall degradation of network performance. The intent of these tests were to identify the channel capacity and degradation in overall performance as more clients requested transmission intervals (uplink). Multiclient downlink testing can identify the effectiveness of the APs service scheduling mechanisms. A secondary interest of this testing is to quantify the fairness, i.e. the percentage of the throughput multiple clients can get in an interval compared to the ideal throughput they could obtain as single clients. Furthermore, it is of interest how this fairness changes as link quality varies from client to client. In baseline multiclient tests all clients are located in close proximity so they can maintain strong links. The setup is shown in Figure 5.27a. The results of this test are then compared to a 'diverse-link' test whereby clients are spaced such that they obtain substantially different SNRs and data-rates. This setup is pictured in Figure 5.27b. Fairness degradation for weak-link clients would indicate that the service period allocation is based off of time intervals, which the high data-rate clients can transmit more information due to use of higher data rates.

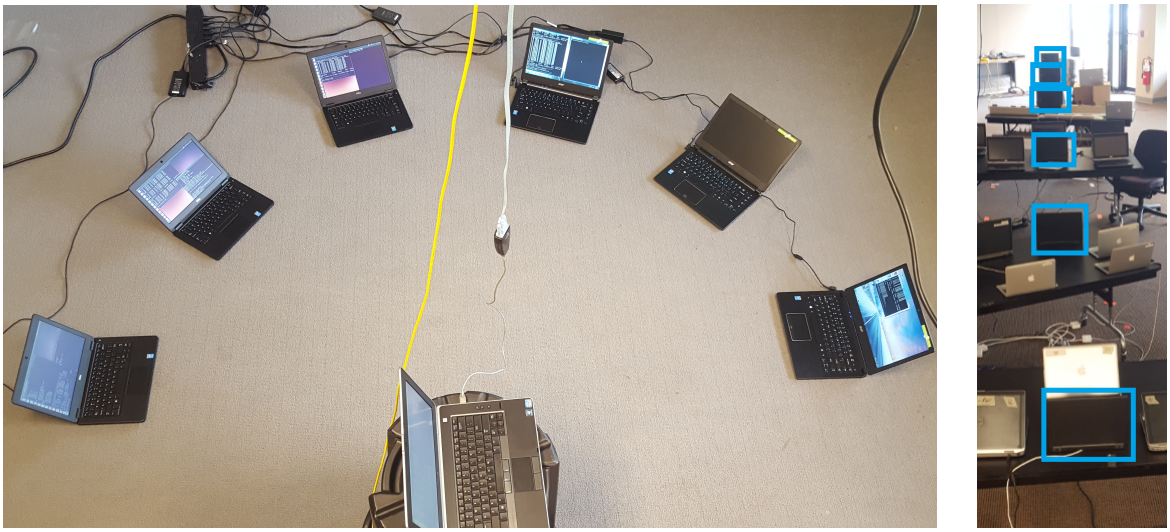

(a) Strong Link (b) Diverse Link

Figure 5.27: Multiclient STA Placement

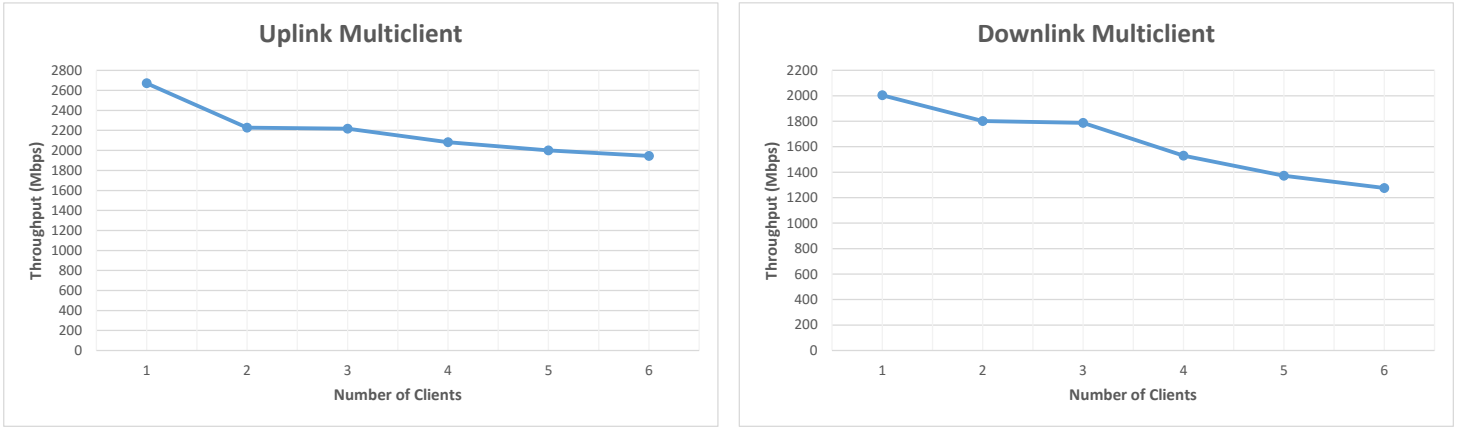

Figure 5.28: Strong Link Multiclient

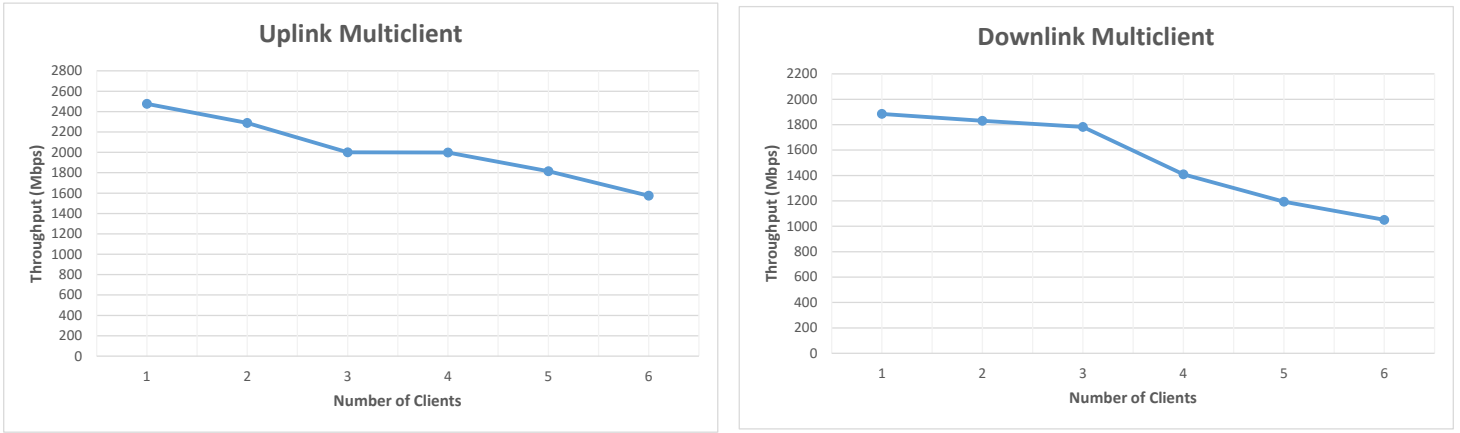

Figure 5.29: Diverse Link Multiclient

The multiclient test results presented in Figures 5.28 and 5.29 show many similarities. As expected the cumulative performance decreases as the network scales to add more stations. Due to the very high throughput demanded by WiGig, it is likely on the uplink that collisions occur frequently. Since the testing conducted was UDP, there are no ACK packets other than the link-layer acknowledgments. This means that collisions are much less likely to occur on downlink, and the fairness allocation is essentially determined by the scheduling mechanisms on the AP and the CPU affinity of each of the performance test processes. In these tests, all Iperf processes were invoked with the same priority so as expected there is a very equal division of throughput amongst clients.

As expected, the strong link tests have higher cumulative performance compared to the weak link tests. In the weak link tests, several of the further away clients cannot obtain the higher MCS rates. Hence, whenever they are granted transmit opportunities they pass less data in a fixed timing interval relative to a strong link client which locks a higher PHY rate. For uplink, the throughput drops from over 2.65 Gbps to 2 Gbps from one to six clients in the strong link tests. In the diverse link, the throughput drops from 2.5 Gbps to 1.6 Gbps. This indicates a loss of 25% and 36% respectively. The higher drop in capacity in the diverse link tests is to be expected.

Another interesting point is that the Dell Linux-based clients performed better than the Acer Windows-based clients. According to the vendor they have been able to achieve higher performance on the Dell laptops due to architectural enhancement that can be enabled in the BIOS. It is also expected that the Linux networking stack is more suitable for multi-Gbps performance, in comparison with the Windows stack. For whatever reason(s), the Windows laptops single-client performance caps out around MCS8 @ 2Gbps. For this reason, the Linux clients were always the first, second and third clients added to the multiclient tests, and the Linux clients were the fourth, fifth and sixth. This was done so that in the single-client baseline, the performance was not limited by the Acer clients bottleneck. Similarly in the diverse link tests the Acer clients were placed furthest away (32 ft, 40 ft, 48 ft) for the same reasoning.

Chapter 6

Conclusion

The first aim of this thesis was to evaluate the efficacy of 60 GHz WiGig (IEEE 802.11ad) to augment wireless networks and enhance performance in real-world indoor environments such as long halls, cubical farms, open office spaces and conference rooms. A secondary objective was to engineer a custom PoC platform to investigate the effects of antenna placement, orientation and topology on the coverage area provided by a 60 GHz radio. An ancillary objective of the PoC was to identify the system engineering design challenges associated with supporting this new WLAN amendment.

The results of this research show that early generation WiGig platforms display very promising performance and reliability in typical office environments. For long hall LOS testing, the PoC platform outperformed theoretically expected results which neglected the negative effects of material absorption, scattering and multipath. Both the client-style and AP-style antenna topologies tested could maintain application layer TP as high as 1.8 Gbps up to 160 ft, which is significantly further than expected in other studies. This indicates high feasibility for WiGig links to be used as wireless backhaul with easier setup than cabled alternatives and higher performance than 1 Gbps Ethernet. The continued operation of the client-style antenna in the modified COTS laptops indicates that the power consumption and heat generated is not an insurmountable barrier for integration into portable devices.

The performance shown by Intel's docking hardware show promising results for removing wires in the PAN and WiDi or other WBE functionality. The hardware showed bottlenecks around 1.3 Gbps indicating that the 60 GHz link could carry Ethernet, a compressed display format, and standard peripherals. As the tech becomes more refined it is likely that the bottleneck will be removed and the dock will support the higher TP single

carrier rates (up to 4.3 Gbps PHY rate). This could unlock uncompressed display and high throughput wireless storage expansion for example.

The environmental testing showed high applicability of WiGig in all sorts of office environments. As expected, the small meeting room could leverage multi-Gbps performance at all points which makes it very suitable for teleconferencing, and media center applications. In the presence of obstructions, the short distance to reflection surfaces such as walls and desks allowed the link to quickly adjust to a more robust MCS while maintaining Gbps performance which is still higher than cutting-edge Wi-Fi. The 60 GHz service area extended outside of the meeting room by using the open door as a first order reflection, and also service could be maintained through the glass windows without noticeable degradation in performance.

In the open office space the performance of WiGig was drastically higher than that offered by state of the art 802.11ac hardware. Even at the very edges of the room at high angles from the bore-sight direction of the AP antenna, hundreds of Mbps could be maintained. When blockages were introduced in the open office, the link was less able to recover as there were not many close reflection paths. However in a realistic open office, there may be more pillars or other objects which can be used to mitigate obstruction degradation. Otherwise, FST offers a promising way to mitigate interruption and loss of service in these extreme scenarios.

The results of the cubical farm testing showed that multi-Gbps performance was attainable in nearly all of the NLOS cubicles even at far distances. In the NLOS environment however, the importance of client station placement and orientation was much more noticeable. Within a given cubicle the variation in attainable performance was drastic depending on the client angle, elevation and location. The position of these 'hotspots' and 'shadowed zones' was fairly intuitive, indicating that quasi-optical propagation models for 60 GHz waves are a reasonable approximation. When the station approached the edge of the cubical walls, the performance would increase until it would eventually behave as a LOS link, suggesting that link could be maintained through the use of both reflection and diffraction. With the use of strategically located wireless relays and repeaters, or even detachable external client antennas, the performance seen in this sort of environment could be nearly identical to that of a LOS environment like the meeting room. An example of an intelligent antenna placement would be every several meters at the at the top of cubicle boundaries. The high frequency re-use and channel availability at 60 GHz could even allow these repeaters and relays to echo the signal on a different channel (or on the same channel with spatial sharing coordination) without decreasing the wireless bandwidth as is the caveat with conventional Wi-Fi repeaters.

The open office antenna characterization testing also displayed some important results with regards to antenna orientation, placement, and topology. If a highly directional antenna is used, it should either be used in a diversity array to minimize shadowed zones, or it should be placed at the edge of the open space with its maximum EIRP direction parallel with the ground plane so that no stations are affected by the shadowing. Also, if the 60 GHz EM is side-car mounted, there is a noticeable benefit to placing the antenna module slightly protruding downward (i.e. periscoped) so that the AP housing does not create unnecessary shadowed regions. The use of a low-loss extension cable could also prevent such shadowing. The unhoused vs housed tests displayed a minor variation in performance indicating that the iota of attenuation added by the thin plastic housing is not significant enough to justify leaving the module exposed.

The obstruction testing attempted to qualify attainable performance in dynamic environments with transitory blockages. When human blockages enter the link service quickly downgrades to lower rates. In the small conference room, nearby reflection surfaces allowed it to quickly rebound to lower MCS rates while maintaining slighly over one Gbps of wireless performance. In the open office, there were no nearby reflection surfaces causing more significant degradation. Even though the link was interrupted, service was maintained at the lowest MCS which still offered 200 Mbps, comparably higher than typical 2.4/5 GHz configurations. When the link obstruction was removed, the link would quickly recover to the optimal MCS within a fraction of a second. In the persistent obstruction tests there was a non-negligible amount of variation in performance and the link did not approach a steady state TP. Overall, human obstructions can cause the link to degrade but nearby surfaces can prevent significant loss of performance, and even in non-ideal environments service is not completely interrupted. The wooden door obstruction tests showed that 802.11ad hardware is capable of maintaining a DMG link of approx 2 Gbps through a thin wooden door. However qualitatively, it was difficult for the station to discover the AP through the door if the link was not already established. Futhermore, during the wooden door obstruction tests the endpoints were less than ten feet away. At higher distances, it is possible service would be interrupted. Lastly, glass windows did not seem to affect practical performance of the 60 GHz links given that higher order MCS rates are already bottlenecked by system side processing limitations anyway. Materials with low absorption coefficients at 60 GHz such as glass windows and ceiling tiles do slightly degrade the signal which can cause the range of higher order MCS rates to be lower, but during static tests these effects are less apparent.

The orientation testing showed that station placement can be an important decision in static deployments such as home media centers, or sync-and-go kiosks. If WiGig is to be adopted into portable devices such as laptops and smartphones, consumers should be

aware of where the antenna modules reside so that they can understand how the device should be oriented in order to leverage the maximum benefit of this 60 GHz technology. The testing revealed that while the effects are minor at lower distances (e.g 8 ft), at larger runs (e.g 24 ft), the degradation due to antenna misalignment is more pronounced. When comparing the orientation testing vs antenna topology, the use of extra transmit chains and combining horizontally and vertically polarized antenna elements (XPR), the degradation effects of antenna misalignment can be reduced to a certain extent.

The mobility tests attempted to measure the performance degradation due to station mobility. From the lowest speed to highest speed tests it was apparent that increased station velocity caused higher variance (i.e. lower reliability) and lower average performance, but not by a drastic amount. Even at a fast jogging speed, multi-Gbps performance could be maintained indicating that the beamforming configuration can be adjusted in such quick intervals that station mobility is not an insurmountable barrier at 60 GHz.

The spatial re-use testing showed no measurable adjacent channel performance degradation due to the spectral mask used in the standard. Due to high channel availability in most geographical regions RRM is a non-issue for adjacent channel networks. By using multiple channels systems could achieve aggregate wireless TP up to tens of Gbps even just using single carrier rates. The spatial re-use testing showed that co-channel networks do interfere at a link spacing of 16 feet which indicates that the beams are not extremely narrow. At 16 foot spacing the spatial re-use factor β_{avg} was around 1.33. The CCI increased fairly significantly at 8 feet ($\beta_{avg} = 1.08$) but even links which coincided could co-exist with a spatial re-use factor approaching one ($\beta_{avg} = 1.01$). The PoC platform displayed service areas covering the entire open office indicating that in both these tests the network overlap was nearly 100%. Despite the noticeable CCI degradation, the tested distances (16 ft, 8 ft, and cross-link) are very extreme and typical networks could be deployed with much higher separation or even no overlap in service area. Furthermore, both links were being tested to their peak TP which greatly increases contention and packet loss. Realistic deployments would have more silent periods and lower channel utilization making collisions less likely. The peak performance and spatial re-use factors measured for the Cisco EM were higher compared to those of the reference designs tested the summer prior. This is atypical, since higher TP would result in more collisions lowering frequency re-use. The twofold improvement could be due to driver enhancements indicating there is still room for implementation improvement and refinement. The next generation of 60 GHz WLAN (802.11ay) allows channel bonding and advanced features such as MU-MIMO which will further increase wireless capacity and reliability at 60 GHz.

The multiclient testing showed that there is a measurable decrease in cumulative performance and spectral efficiency as more contending stations enter the network. This is a

reality due to the backoff mechanisms associated with the multiple access protocol used by 802.11ad. Even with up to 6 stations, the aggregate performance was around 2 Gbps showing that WiGig is capable of delivering VHT to multiple endpoints simultaneously. Each of the contending clients were requesting the maximum bandwidth they could obtain, which exacerbates the degradation caused by collisions. A more realistic load profile likely would contain more temporal variation in each endpoints requested TP which could reduce collisions and backoff frequency. The testing conducted in this research aimed to quantify an upper-limit for multiclient performance degradation. As demonstrated by the strong link tests, endpoints with similar SNRs evenly divided the available TP indicating the fairness of the random exponential backoff. In the weak link tests, distant clients with decreased SNRs were limited to lower MCS rates resulting in lower average performance in a given access interval, when compared to nearby clients. Since the total channel capacity was measured as that of a strong-link client, the diverse link tests displayed multiclient degradation up to 35%, compared to the 25% measured in the strong link tests. Realistically, the diverse link 'capacity' could be considered to be the average of each stations peak performance since the distant clients negatively skew the aggregated TP.

In conclusion, 802.11ad is a very promising augmentation to the WLAN family and will inevitably provide the next enhancement in wireless capacity for home and business enterprises. As the technology becomes more understood, unique applications will continue to emerge where fiber is infeasible and high performance is required. The PHY and MAC protocols in the standard provide a high amount of versatility making WiGig suitable in even dynamic environments with obstructions and mobile stations. In static LOS, 802.11ad is a cheap and effective means to obtain DMG. Compatibility with existing WLAN standards will allow quicker client and infrastructure adoption due to Wi-Fi's ubiquity and popularity amongst consumers. Despite small variations in peak performance increased TP will increase application reliability. By setting application demands lower than the peak TP by a small safety margin and with the use of larger application buffers, 60 GHz DMG can appear as reliable as wired alternatives. The massive practical performance boost enables use of lower compression ratios for WiDi and other wireless transfers. This reduces de/compression latency allowing a reduction in end-to-end delay for large transfers. High frequency reuse at 60 GHz allows for increased physical security, high network density, no ACI, and even co-channel coexistence and spatial re-use with smart AP co-ordination. Lastly, adding 60 GHz to WLAN deployments will increase performance of 2.4/5 GHz clients by virtue of offloading the most bandwidth hungry small-cell applications. This allows less congestion and contention on the lower bands making access intervals more abundant to service clients which are not 60 GHz capable (either due to placement or lack of technology support). As an extremely new technology there is a need for continued

research into 60 GHz WiGig.

6.1 Future Work

Due to the high frequency and high data bandwidth of these systems there needs to be a significant research and design effort into efficient high performance 60 GHz RF hardware components such as antennas, mixers and filters. The competing requirements for high performance and low power consumption creates unique challenges for successful implementation in silicon and ceramic based substrates. Enhancements to the hardware components and fabrications technologies will enhance the effectiveness of this critical wireless technology.

Furthermore, there are still many features of the standard yet to be supported by early technology vendors: For example FST with support for seamless tri-band integration, or support for more than eight clients. Once these features are enabled, doors open for even more advanced research into RRM, minimization of degradation due to roaming and network scaling etc. As the software drivers for 60 GHz hardware are refined more controllability may be exposed to the user through application programming interfaces. With fine-tuned control over beamforming, MCS rates, or link-level debug information (e.g. RSSI and packet-loss statistics), investigating the effects of minor changes and variations is more feasible.

Lastly, accurately modeling the 60 GHz channel poses a significant challenge and commercial hardware inexpensive and difficult to obtain. Wireless instrumentation for traditional Wi-Fi is cannot be used on WiGig due to lack of tuned frequency support and bandwidth limitations. Research efforts into sounding 60 GHz channels can provide valuable insight into the effects of environment, local scatterers, obstructions and obstacles, and mobility. Enhancement of computer simulation models will allow for more fine tuned simulation of 60 GHz performance which could lead to less discrepancy between expected and realized throughput and reliability.

References

[1] John M. Shea. History of Wireless Communication. http://wireless.ece.ufl.edu/jshea/HistoryOfWirelessCommunication.html. Accessed: 2016-10-23.

[2] Wi-Fi Alliance: Organization official industry association website. http://www.wi-fi.org/who-we-are. Accessed: 2016-10-23.

[3] Wi-Fi Alliance. Wi-Fi®device shipments to surpass 15 billion by the end of 2016. http://www.wi-fi.org/news-events/newsroom/wi-fi-device-shipments-to-surpass-15-billion-by-end-of-2016. [Online; posted July 29th, 2016] Accessed: 2016-11-08.

[4] Machina Research. Global Machine to Machine Forecast. https://machinaresearch.com/news/global-m2m-market-to-grow-to-27-billion-devices-generating-usd16-trillion-revenue-in-2024/. Accessed: 2016-11-08.

[5] Ubee Interactive. Understanding Technology Options for Deploying Wi-Fi - How Wi-Fi Standards Influence Objectives. http://www.ubeeinteractive.com/sites/default/files/Understanding%20technology%20options%20for%20deploying%20WiFi%20.pdf. Accessed: 2016-10-23.

[6] Rebel Roam. Rebel Roam Mobile Data Consumption Growth. http://www.rebelroam.com/2016/07/29/mobile-internet-data-consumption-growing-50-per-year-coach-wifi/. [Online; posted July 29th, 2016] Accessed: 2016-11-08.

[7] Rajesh K. Gigabit Wi-Fi - 802.11ac Advantages & Limitations. http://www.excitingip.com/3280/gigabit-wi-fi-802-11ac-advantages-limitations/. Accessed: 2016-10-24.

[8] M. Zaaimia, R. Touhami, A. Hamza, and M. C. E. Yagoub. Design and performance evaluation of 802.11ad PHYs in 60 GHz multipath fading channel. In *2013 8th International Workshop on Systems, Signal Processing and their Applications (WoSSPA)*, pages 521–525, May 2013.

[9] King, Kimberley Brynn. Digitally-Assisted RF-Analog Self Interference Cancellation for Wideband Full-Duplex Radios. UWSpace, 2016.

[10] Yun Liao, Kaigui Bian, Lingyang Song, and Zhu Han. Poster: Full-duplex WiFi: Achieving Simultaneous Sensing and Transmission for Future Wireless Networks. In *Proceedings of the 16th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, MobiHoc '15, pages 381–382, New York, NY, USA, 2015. ACM.

[11] Sachin Katti. Full Duplex Radios: From Impossibility to Practice. In *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies*, CoNEXT '13, pages 301–302, New York, NY, USA, 2013. ACM.

[12] Andrei Stroe. Wi-Fi Hidden Station Problem, 2009. File: Wifi hidden station problem.svg.

[13] James F. Kurose and Keith W. Ross. *Computer Networking: A Top-Down Approach (4th Edition)*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2007.

[14] L. Lu, X. Zhang, R. Funada, C. S. Sum, and H. Harada. Selection of modulation and coding schemes of single carrier PHY for 802.11ad multi-gigabit mmWave WLAN systems. In *2011 IEEE Symposium on Computers and Communications (ISCC)*, pages 348–352, June 2011.

[15] Katherine Finnell and Margaret Rouse. Millimetre Wave (MM wave) definition. http://searchtelecom.techtarget.com/definition/millimeter-wave-MM-wave. Posted April 2015.

[16] Federal Communications Commission. FCC Online Table of Frequency Allocations. https://transition.fcc.gov/oet/spectrum/table/fcctable.pdf. Revised August 31, 2016.

[17] W. Haiming, H. Wei, C. Jixin, S. Bo, and P. Xiaoming. IEEE 802.11aj (45 GHz): A new very high throughput millimeter-wave WLAN system. *China Communications*, 11(6):51–62, June 2014.

[18] Agilent. Wireless LAN at 60 GHz - IEEE 802.11ad Explained - *Application Note*. http://cp.literature.agilent.com/litweb/pdf/5990-9697EN.pdf. Accessed: 2016-05-06.

[19] Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification - Amendment: Enhanced Throughput for Operation in License-Exempt Bands Above 45 GHz. IEEE 802.11ay, 2016. Piscataway, NJ.

[20] T. Nitsche, C. Cordeiro, A. B. Flores, E. W. Knightly, E. Perahia, and J. C. Widmer. IEEE 802.11ad: directional 60 GHz communication for multi-Gigabit-per-second Wi-Fi [Invited Paper]. *IEEE Communications Magazine*, 52(12):132–141, December 2014.

[21] S. Shankar N., D. Dash, H. El Madi, and G. Gopalakrishnan. WiGig and IEEE 802.11ad - For multi-gigabyte-per-second WPAN and WLAN. *ArXiv e-prints*, November 2012.

[22] David A. McGrew and John Viega. The Security and Performance of the Galois/Counter Mode (GCM) of Operation. In *In INDOCRYPT, volume 3348 of LNCS*, pages 343–355. Springer, 2004.

[23] Ngan Tengyuen. Router Range and Distance Comparison. https://www.geckoandfly.com/10041/wireless-wifi-802-11-abgn-router-range-and-distance-comparison/. Posted October 11th, 2016.

[24] Mehmet R. Yuce. *Ultra-Wideband and 60 Ghz Communications for Biomedical Applications (1st Edition)*. Springer, New York, US, 2014.

[25] Matthew Gast. *802.11ac: A Survival Guide*. O'Reilly Media, Inc., 1st edition, 2013.

[26] TP Link. Talon AD7200 Multi-Band Wi-Fi Router - Specifications. http://www.tp-link.com/en/products/details/cat-9_AD7200.html. Released January 2016.

[27] Netgear. Nighthawk X10 Smart Wi-Fi Router - Specifications. https://www.netgear.com/landings/ad7200/. Released October 2016.

[28] Acelink. Br-6774AD Tri-Band Wi-Fi Router - Specifications. http://www.acelink.com.tw/BR-6774AD.html. Released November 2016.

[29] Accton. Metrolink 60 GHz Client Module. https://fccid.io/HED-ML60MDSB. FCC Certified September 2016.

[30] Peraso. Client 802.11ad USB Dongle. http://www.perasotech.com/consumer/. In development.

[31] Intel. Intel Tri-Band Wireless AC 17265 Wi-Fi Router - Specifications. http://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/tri-band-wireless-ac17265-brief.pdf. Released October 2016.

[32] Tai Huang. Feasibility Study of IEEE 802.11ad for Vehicle-to-X Communication. *Chalmers University of Technology*, 2015.

[33] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parmesh Ramanathan. 60 GHz Indoor Networking Through Flexible Beams: A Link-Level Profiling. *SIGMETRICS Perform. Eval. Rev.*, 43(1):71–84, June 2015.

[34] Daniel Halperin, Srikanth Kandula, Jitendra Padhye, Paramvir Bahl, and David Wetherall. Augmenting Data Center Networks with Multi-gigabit Wireless Links. *SIGCOMM Comput. Commun. Rev.*, 41(4):38–49, August 2011.

[35] Hany Assasa and Joerg Widmer. Implementation and Evaluation of a WLAN IEEE 802.11ad Model in ns-3. In *Proceedings of the Workshop on Ns-3*, WNS3 '16, pages 57–64, New York, NY, USA, 2016. ACM.

[36] Anique Akhtar and Sinem Coleri Ergen. Efficient Network Level Beamforming Training for IEEE 802.11ad WLANs. In *Proceedings of the International Symposium on Performance Evaluation of Computer and Telecommunication Systems*, Spects '15, pages 1–6, San Diego, CA, USA, 2015. Society for Computer Simulation International.

[37] Carlos Cordeiro, Dmitry Akhmetov, and Minyoung Park. IEEE 802.11ad: Introduction and Performance Evaluation of the First Multi-Gbps Wi-Fi Technology. In *Proceedings of the 2010 ACM International Workshop on mmWave Communications: From Circuits to Networks*, mmCom '10, pages 3–8, New York, NY, USA, 2010. ACM.

[38] Nathan Ida. *Engineering Electromagnetics*. Springer Publishing Company, Incorporated, 2nd edition, 2007.

[39] H. T. Friis. A Note on a Simple Transmission Formula. *Proceedings of the IRE*, 34(5):254–256, May 1946.

[40] C. Cordeiro. Evaluation of Medium Access Technologies for Next Generation Millimeter-Wave WLAN and WPAN. In *2009 IEEE International Conference on Communications Workshops*, pages 1–5, June 2009.

[41] E. Perahia, C. Cordeiro, M. Park, and L. L. Yang. IEEE 802.11ad: Defining the Next Generation Multi-Gbps Wi-Fi. In *2010 7th IEEE Consumer Communications and Networking Conference*, pages 1–5, Jan 2010.

[42] Radiometry and the Friis Transmission Equation. *American Journal of Physics*, 81, 2012.

[43] Jason D. McKinney V. J. Urick, Keith J. Williams. *Fundamentals of Microwave Photonics*. Wiley, 1st edition, 2015.

[44] Federal Communications Commission. Millimeter Wave Propagation: Spectrum Management Implications. In *FCC Bulletin 70A*, page 12. July 1997.

[45] Shigeaki Hakusui. Fixed wireless communications at 60 ghz unique oxygen absorption properties, 2001.

[46] J. Lu, D. Steinbach, P. Cabrol, P. Pietraski, and R. V. Pragada. Propagation characterization of an office building in the 60 GHz band. In *The 8th European Conference on Antennas and Propagation (EuCAP 2014)*, pages 809–813, April 2014.

[47] IEEE. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification. *IEEE 802.11*, (802.11), 1997.

[48] Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification - Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band, 2012.

[49] Dinesh Bharadia, Emily McMilin, and Sachin Katti. Full Duplex Radios. *SIGCOMM Comput. Commun. Rev.*, 43(4):375–386, August 2013.

[50] Leon W. Couch, II. *Digital and Analog Communication Systems*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 3rd edition, 1990.

[51] Saeed Ur Rehman, Stevan Berber, Akshya Swain, and Wayne Holmes. Modeling the Impact of Deferred Transmission in CSMA/CA Algorithm of IEEE 802.15.4 for Acknowledged and Unacknowledged Traffic. In *Proceedings of the 8th ACM Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks*, PE-WASUN '11, pages 25–32, New York, NY, USA, 2011. ACM.