

60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling

Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, Parmesh Ramanathan

{sur2, vvenkateswar}@wisc.edu, {xyzhang, parmesh}@ece.wisc.edu

University of Wisconsin-Madison, USA

ABSTRACT

60 GHz technology holds tremendous potential to upgrade wireless link throughput to Gbps level. To overcome inherent vulnerability to attenuation, 60 GHz radios communicate by forming highly-directional electronically-steerable beams. Standards like IEEE 802.11ad have tailored MAC/PHY protocols to such flexible-beam 60 GHz networks. However, lack of a reconfigurable platform has thwarted a realistic proof-of-concept evaluation. In this paper, we conduct an in-depth measurement of indoor 60 GHz networks using a first-of-its-kind software-radio platform. Our measurement focuses on the link-level behavior with three major perspectives: (i) coverage and bit-rate of a single link, and implications for 60 GHz MIMO; (ii) impact of beam-steering on network performance, particularly under human blockage and device mobility; (iii) spatial reuse between flexible beams. Our study dispels some common myths, and reveals key challenges in maintaining robust flexible-beam connection. We propose new principles that can tackle such challenges based on unique properties of 60 GHz channel and cognitive capability of 60 GHz links.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design - Wireless communication;

C.4 [Performance of Systems]: Measurement techniques

Keywords

60 GHz, millimeter-wave, IEEE 802.11ad, software-radio

1. INTRODUCTION

Driven by the growing population of mobile devices and bandwidth-hungry applications, mobile network traffic is set to explode in our near future. Industry research predicts that aggregate wireless traffic will increase by 1000× within the next decade [1]. Legacy WiFi and cellular networks fall

short of such demand because of the scarce spectrum resources in the underlying microwave band, despite tremendous research efforts in improving spectrum efficiency. Moving to higher frequencies, however, copious spectrum exists. Most remarkably, the 57 GHz – 64 GHz spectrum, colloquially known as the 60 GHz millimeter-wave (mmWave) band, allows for unlicensed access across majority of the world. Two IEEE 60 GHz MAC/PHY standards, 802.11ad [2] and 802.15.3c [3], have recently been ratified for local and personal area networking. With up to 7 Gbps of data rate, they are ready to support next-generation wireless applications, *e.g.*, uncompressed video streaming, cordless computing, and even wireless data centers [4]. A similar paradigm has been advocated by industry to realize 5G cellular networks [5].

Nonetheless, the eventual success of such conceptual development depends on system designs that account for the unique characteristics of 60 GHz networks. In particular, 60 GHz links are highly vulnerable to attenuation. The free-space path loss scales as $1/\lambda^2$, where λ is the carrier wavelength [6]. So, 60 GHz is 21.6 dB worse than 5 GHz and 28 dB worse than 2.4 GHz that is being used for current WiFi. On the other hand, 60 GHz links can be made highly directional by compacting many miniature antenna elements into a *phased-array*. Given the same form-factor, directivity gain scales as $1/\lambda^2$, which can compensate the path loss drawback. In IEEE 802.11ad, for example, a link quality gain of 20 dB can be easily achieved using a 16-element phased-array [7]. However, highly directional links induce a new sets of challenges, particularly due to human blockage and device mobility. Although phased-arrays are envisioned to overcome the hindrance by electronically steering beams, little experimental research has been conducted to validate the efficacy.

To date, evaluation of 60 GHz protocols mostly relies on simulations complemented by analytical/empirical propagation models [6]. However, given the diverse beam patterns and sophisticated multipath reflections, it is virtually infeasible to faithfully reproduce mmWave channel profiles with a unified model. This problem becomes especially pronounced in dynamic environment with human blockage, device mobility, and interference between directional beams. Recent commercial-off-the-shelf (COTS) 60 GHz devices [8, 9] enabled transport and application layer throughput profiling. However, lack of access to their propriety MAC/PHY implementation has left many open conjectures *w.r.t.* how 60 GHz networks behave especially under environment dynamics.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGMETRICS'15, June 15–19, 2015, Portland, OR, USA.

Copyright © 2015 ACM 978-1-4503-3486-0/15/06 ...\$15.00.

<http://dx.doi.org/10.1145/2745844.2745858>.

In this paper, we conduct a microscopic measurement of 802.11ad-based 60 GHz indoor WLAN, with an emphasis on how the flexible beams impact link-layer performance, and how they react under blockage/mobility. To this end, we build a first-of-its-kind 60 GHz software-radio testbed, called *WiMi*. A WiMi transmitter can reconfigure its carrier frequency, output power and waveform, which can be detected and processed by a receiver in real-time. Link-level statistics, such as Received Signal Strength (RSS), can be monitored in fine-resolution (≤ 2 ms) even without higher-layer association between transmitter and receiver (unlike COTS devices). WiMi uses mechanically steerable horn antennas to emulate 802.11ad phased-array, and allows plug-and-play installation of antennas with different beamwidths.

Profiling Link-level Characteristics. We leverage WiMi’s reconfigurability to examine 60 GHz flexible-beam links in indoor environment from three perspectives:

(i) *Single static link.* We identify key factors that determine the attenuation models a 60 GHz link should follow in indoor environment. Such factors explain the inconsistencies in prior channel measurement studies that use channel sounder hardware [10, 11] or COTS devices [4]. We then investigate the feasibility of establishing a mmWave MIMO link, by transforming WiMi into a virtual MIMO platform. We further analyze how the link rate/range scales with narrower beamwidth under regulatory power constraints.

(ii) *Link behavior under environment dynamics.* 802.11ad transmitters and receivers rely heavily on a beam searching protocol that aligns their beam directions to maximize link SNR, or detours obstacles via reflection paths. We examine the efficacy and cost of beam searching, along with other pathologies that may eventually nullify the benefits of flexible beams. In addition, we identify design principles that can minimize the impact of link outage caused by human blockage or device motion.

(iii) *Multi-link spatial reuse.* Indoor 60 GHz beams differ from the ideal fan-shaped patterns in theoretical models. We measure the impacts of such imperfectness on spatial reuse. Further, we investigate the impact of environment dynamics on 802.11ad-like MAC, which realize spatial reuse by establishing a conflict graph between directional beams.

Summary of insights from measurement study. Our measurement study leads to several interesting sets of insights, which we summarize as follows:

(i) Existing propagation models typically abstract narrow-beam advantage into an antenna gain factor. We found that beamwidth, along with factors like antenna height, also determine how closely the signal attenuation follows theoretical models. Unlike conventional perception that 60 GHz beams behave in a pseudo-optical manner, we found highly directional beams suffer less penetration loss across typical obstacles (except human body) in an office environment, and coverage can be achieved beyond a single room. A side discovery is that 60 GHz MIMO gain becomes correlated with link distance, instead of antenna distance as assumed in communications-theoretic models [6].

(ii) In device motion scenarios, we found the 802.11ad’s beam searching algorithm costs more channel time than the Gbps data transmission, and thus degrades throughput even for a 22.5° beam with a relatively small search space. 802.11ad proposes a quasi-omni mode aiming to accelerate the AP-client discovery procedure by widening beams, yet we found this rarely helps because the 60 GHz channel itself exhibits a

densely concentrated Angle-Of-Arrival (AOA) pattern. We also discover, for the first time, that 802.11ad links themselves tend to be asymmetric due to sophisticated interaction between beam patterns and environment dynamics.

In human blockage scenario, we found the effectiveness of beam searching is sensitive to blockage position and reflectivity of environment. From both scenarios, we derive and validate two principles that can combat link dynamics: (a) sparsity of AOA can be leveraged to reduce beam searching space and thus overhead; (b) sensitivity of 60 GHz links themselves can be exploited to detect blockage/motion and tackle link outage in a proactive way.

(iii) Due to reflection and strong antenna side lobes, even super-narrow beams (*e.g.* 3.4°) can leak signals causing interference and degrade spatial reuse. Conflict-graph based approach, albeit effective in static scenarios [12], can significantly disturb network performance due to a tradeoff between responsiveness and overhead. It is more sensible to detect device motion, and isolate mobile nodes by allocating separate TDMA slots.

Our Contributions. General properties of 60 GHz channel, such as directionality and vulnerability to blockage, are already well known in communications models [10, 11, 13]. The new contributions of this study lie in its microscopic perspectives on *network* performance, and particularly in: (i) novel link-layer measurement methodologies based on a 60 GHz software-radio platform; (ii) new measurement observations on the link-level behavior of flexible, directional 60 GHz beams under environment dynamics; (iii) new design principles to realize efficient beam searching and adaptation to blockage/mobility, harnessing the inherent sparsity and sensitivity of 60 GHz channels. The data set from this measurement study will be released to the research community through our project repository [14]. The empirical traces can drive the simulation of design choices such as network architecture, scheduling mechanisms and adaptation algorithms to combat blockage/mobility.

2. BACKGROUND AND METHODOLOGY

2.1 IEEE 802.11ad Standard for 60 GHz WLANs

An 802.11ad WLAN comprises multiple stations (STAs) that can establish peer-to-peer (P2P) mode connection between each other, or infrastructure mode connection with an access point (AP). Each node executes a beamforming and beam-steering mechanism to form directional transmit/receive beams using a phased-array antenna. The AP maintains a hybrid MAC scheduler to centrally coordinate all the directional links in the network.

Beamforming and Beam-steering: A phased-array antenna comprises multiple omni-directional antenna elements that together form a radio-frequency “lens”. Beamforming is equivalent to reshaping the lens to form different beam patterns (widths and directions). More specifically, each beam pattern is generated by assigning a vector of phase-shifts to the antenna elements. Given M elements, a discrete set of $2M$ beam patterns can be generated by an 802.11ad phased-array [2]. The actual beamwidth and direction depends on the phase-shift values, and the narrowest beamwidth depends on how large M is.

802.11ad runs an iterative beam searching protocol to align the transmitter and receiver’s beams (details in Section 4.1). Besides, a *quasi-omni* beam pattern (Figure 1), which




Figure 1: Fine beam and its quasi-omni beam.




Figure 2: IEEE 802.11ad beacon interval structure.

is relatively wider, is used opportunistically to improve the rendezvous probability between AP and STAs’ beams. The antenna gain of a quasi-omni beam can be at most 15 dB lower than that of the narrowest beam [2].

Hybrid MAC Layer: 802.11ad uses a hybrid MAC to coordinate the directional communication between multiple devices. Simply put, in infrastructure mode, the AP can enforce a *deterministic TDMA schedule* (a.k.a. Service Period) to communicate with multiple STAs, whose beam directions are trained periodically at the beginning of a beacon interval. In P2P mode, a *semi-contention scheduler* is used, where a STA can contend to reserve channel by sending a RTS request to the AP. The AP confirms the reservation with a CTS. The AP learns the interference relation between P2P links by collecting periodic interference reports from them (Section 5.2). It uses CTS to permit non-interfering links to transmit concurrently with the one who initiated RTS.

Timing in 802.11ad is based on beacon intervals (Figure 2), separated by directional beacons from the AP at the beginning of the interval (BTI). The AP uses quasi-omni beams to broadcast beacons through all spatial directions. Then, the AP and STAs execute a beam training (searching) procedure in an A-BFT slot, followed by an ATI, where AP announces the schedules for the immediate data transmission interval. The different scheduling modes are allocated in orthogonal time frames but the duration and allocation strategy can be customized and left as implementation dependent.

2.2 Measurement Methodology

We now introduce our 60 GHz software-radio that is built for measurement purpose. The software-radio provides signal-level visibility, and is engineered to profile link-level characteristics of 802.11ad devices using a combination of over-the-air transmission and emulation.

2.2.1 Building the WiMi Software-Radio

Existing COTS 60 GHz platforms (e.g., Wilocity radio [9, 15] and HXI Gigalink radio [4]) can only report higher-layer throughput statistics or long-term average RSS, after a link is established. However, our link-level measurement requires a 60 GHz “sniffer” that can measure fine-grained link statistics even without connecting to the transmitter. This is crucial in studying dynamic scenarios with blockage/mobility, where link outages become the norm. In addition, the measurement platform needs to be programmable *w.r.t.* output power, beam patterns, and signal waveforms. Despite plenty of 2.4/5 GHz reconfigurable radios [16–18],

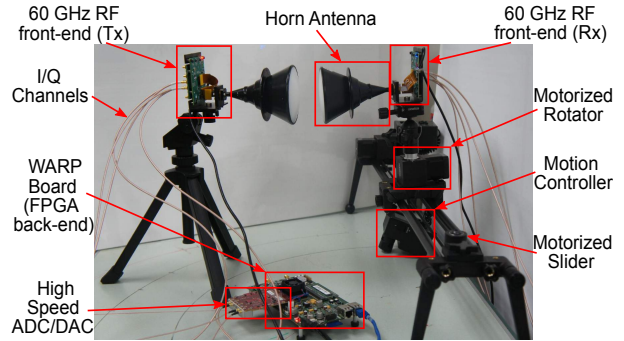


Figure 3: WiMi software-radio platform.

we are unaware of a 60 GHz counterpart. Therefore, we have assembled our own testbed, *WiMi*, to satisfy the measurement requirements. WiMi is a 60 GHz software-radio that leverages the WARP FPGA [16] as baseband signal processing unit, along with a high-speed ADC/DAC module connected to a 60 GHz RF development board [19]. Figure 3 illustrates a pair of WiMi transmitter and receiver and their major building blocks.

2.2.2 Creating and Measuring 60 GHz Links on WiMi

RSS, bit-rate and throughput measurement. Each WiMi node is hosted by a PC that controls the radio hardware and executes MAC/PHY operations through a mix of MATLAB/C drivers. We have developed customized drivers that allow WiMi to send/receive 60 GHz packets at a fine-granularity of 2 ms per packet, each lasting up to 800 μ s. The receiver computes RSS on a per-packet basis, by calculating the average signal power across the 800 μ s. This suffices for evaluating link dynamics caused by environment change.

Due to hardware limitation, WiMi’s bandwidth (245.76 MHz) is incommensurate with 802.11ad radios (2.16 GHz), so it cannot achieve Gbps data-rate. However, the measured RSS and noise floor can be translated into achievable bit-rate following an 802.11ad specific rate table (Table 2 in Appendix C). Such extrapolation approach has been adopted and verified in existing measurement studies [4, 9].

Unless noted otherwise, the transmit power of a WiMi node is always calibrated such that the RSS matches the minimum sensitivity required for the highest data rate, before starting any experiment. The protocol stack on the PC host emulates the 802.11ad MAC components that are pertinent to this study. The emulator uses a virtual clock when enforcing protocol actions including packetization with preambles, beam searching, inter-frame spacing, RTS/CTS, ACK, *etc.* that respect 802.11ad’s default timing parameters. Link throughput is calculated taking into account the achievable bit-rate along with such overhead.

Adapting antenna beam width and directions. Due to the lack of any COTS electronically steerable phased-array antenna at 60 GHz, we emulate 802.11ad’s beam-steering by using a mechanically steerable horn antenna with beamwidth 3.4° and antenna gain of 34 dBi (equivalent to a phased-array with 50×50 antenna elements). Besides, an omni-directional antenna can be plugged in and moved to create a virtual MIMO antenna (Appendix B). We further create a variety of different beam patterns by “reshaping” WiMi’s built-in waveguide module. According to our measurement (Figure 24 in Appendix A), the waveguide pro-




Figure 4: Impact of beamwidth on LOS path loss. Antenna height is 30 cm.




Figure 5: Impact of antenna height on LOS path loss. Beamwidth is 180°.




Figure 6: 2 × 2 MIMO capacity gain for different link distances.

vides approximately 180° beamwidth in horizontal and vertical plane. We leverage mmWave absorbers [20] and create conical shapes opening in front of the waveguide according to the beamwidth requirement. Examples beam patterns generated using this approach are verified in Appendix A. Similar to the horn antenna, beam directions of the emulated beam patterns can be steered mechanically.

3. PROFILING A SINGLE STATIC LINK

We first inspect the coverage and capacity of a standalone 60 GHz link, with a particular emphasis on 802.11ad directional beams.

3.1 Line-of-Sight (LOS) Links

In ideal free-space setting, RF signals’ attenuation loss follows the Friis model (in dB scale):

$$\begin{aligned}
 PL(d) &= 20 \log_{10}\left(\frac{4\pi d}{\lambda}\right) = 20 \log_{10}(4\pi d \cdot \frac{f}{c}) \\
 &= 20 \log_{10}(d) + 20 \log_{10}(f) - 147.56 \quad (1)
 \end{aligned}$$

where d is the link distance, and λ , f , c are the carrier wavelength, signal frequency and light speed, respectively.

Whereas measurement of COTS 60 GHz links showed almost perfect fit to the Friis model [4, 21], PHY-layer channel sounding tests revealed drastic deviations [10, 11]. Owing to the flexibility of WiMi, we have identified two factors that can account for such discrepancies.

Attenuation of different beamwidth. We vary the distance of a LOS link, and measure the propagation loss relative to a reference distance $d_0 = 0.5$ m. The results (Figure 4) highlight two immediate observations: (i) A narrower beam follows the Friis model more closely. (ii) The shorter the distance, the closer the path loss is to the Friis model. In effect, both phenomena may be explained by the multipath reflections that disturb the model at further distances where the wavefront hits obstacles. To corroborate this point, we further vary the antenna position that affects how severe multipath is at certain distance.

Impacts of antenna positioning. We mount a pair of nodes with 180° beamwidth on a variety of heights above the floor. Figure 5 shows that, higher the antennas are, closer the loss is to Friis model, because it takes longer distance before the beams hit the floor to cause multipath reflections. It should be cautioned that the Friis model, being ideal, is not necessarily optimistic. At further distances, measured path loss can be even lower than that in free-space. This is due to a *waveguide effect*, where multipath are chambered to strengthen each other in indoor environment [13].

Implications for 60 GHz MIMO links. Next-generation 60 GHz MIMO network standard is already charted by the IEEE NG60 group [22], aiming to achieve 30+ Gbps rate by

delivering parallel streams of data through multiple omnidirectional antennas simultaneously. Such *multiplexing gain* is achievable only if the antennas experience *uncorrelated* multipath channels [23]. However, the foregoing experiments hint that co-located antennas may experience *similar* Friis loss pattern, thus impairing the MIMO potential. We verify this hypothesis by building a 2 × 2 virtual MIMO on WiMi (see Appendix B), with half-wavelength separation between co-located antennas to minimize coupling effects. We then measure the capacity gain of MIMO over corresponding SISO link, with nodes mounted 30 cm above a table. Theoretically, this setup should achieve close to 2× gain. Yet the results (Figure 6) show the MIMO gain is marginal ($< 1.26\times$) when link distance falls below 1.5 m, and gradually approaches 2 over longer distances. This is again because the parallel MIMO links experience random, uncorrelated multipath only beyond a certain distance.

3.2 Non-Line-of-Sight (NLOS) Links

It is a common perception that 60 GHz signals propagate in a pseudo-light manner and cannot penetrate walls or other obstacles indoor. We now examine this perception in an office environment, with floor plan shown in Figure 7(b).

Blockage loss by common objects indoor. We measure the penetration loss by placing the transmitter and receiver 1 m apart with obstacles in between, and compare with a LOS link of the same distance.

Table 1 shows that, for a 3.4° beam, penetration loss across drywall, wooden door, and cubicle partitions are below 1 dB, leading to only 7.5% of throughput degradation. Both whiteboard and metal sheet (0.1 cm) cause more than 4 dB of RSS loss and 33% of throughput degradation. Given equal beamwidth (30°), 60 GHz suffers more penetration loss than WiFi. However, the loss is significantly lower when using a 3.4° highly directional beam, likely because of less diffusion effects along a concentrated beam. To our knowledge, this phenomenon has not been reported in prior measurement studies that use inflexible channel sounder [10, 11] or COTS devices [8, 9]. As it is easier to build small form-factor directional antennas at 60 GHz [6], this advantage can partly compensate for the vulnerability to penetration loss.

Remarkably, 60 GHz signals experience dramatic attenuation when going through human body or water bottle. When a transmitter sends 3.4° and 30° beams and is fully blocked by a close-by human body, the receiver senses no signal. In contrast, the 30° WiFi link can still maintain connectivity, though at around 43% of throughput loss.

Coverage under penetration loss. Given the surprisingly low penetration loss across typical obstacles, we expect 60 GHz links can cover beyond a single room. For a quantitative validation, we first place a transmitter on a corner, and survey the receiver’s spatial throughput distribution by

Material (thickness)	60 GHz (3.4°)		60 GHz (30°)		WiFi (30°)	
	RSS	T _{put}	RSS	T _{put}	RSS	T _{put}
Drywall (15 cm.)	0.53 dB	7.5%	3.2 dB	23%	1.5 dB	12.5%
Wooden door (4.5 cm.)	0.73 dB	7.5%	4 dB	33.3%	1.2 dB	12.5%
Office cubicle separator (5 cm.)	0.4 dB	7.5%	2.9 dB	23%	0.2 dB	12.5%
Office white board (2.5 cm.)	4.2 dB	33%	7.6 dB	48.7%	2.6 dB	14.3%
Metal sheet (0.1 cm.)	6.5 dB	38.5%	10.1 dB	58.9%	8 dB	25%
Human body	31.2 dB	100%	30.9 dB	100%	12.5 dB	42.9%
Water bottle	29.19 dB	100%	-	-	-	-

Table 1: RSS and throughput loss due to obstacles.

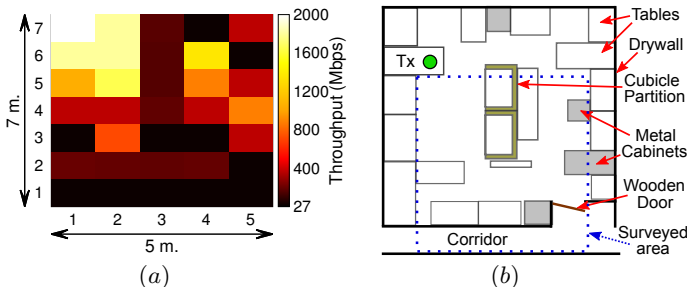


Figure 7: (a) Spatial throughput distribution of a 180° link inside an office and in corridor. Black spots denote grids with lowest throughput. (b) Floor plan of the office.

dividing the measurement region into 1m × 1m grids. Within each grid, we sample the RSS at 4 random points, each lasting 20 s, and calculate the average. Both nodes are configured to a wide beam of 180°, corresponding to the lowest antenna gain and most conservative rate/coverage estimation.

From the result (Figure 7(a)), we see that even the 180° beam can sustain Gbps throughput within 1 to 2 m, and can maintain link connectivity with lowest data rate almost across the entire room and even outside, despite all the obstructions.

To further explore the coverage limit of directional 60 GHz links beyond room level, we configure the nodes to use two different antennas, and measure the rate degradation over distance and across a whiteboard, drywall and wooden door. The transmit power is calibrated, such that a receiver 50 cm away can have minimum RSS required to achieve highest rate. For comparison, we also calibrate a 2.4 GHz WiFi link such that the RSS at 50 cm is 28 dB higher than 60 GHz (to account for the Friis loss, see Section 3.1). Figure 8 shows that an omni-directional link quickly reaches a coverage limit of 5 m, especially after wall obstruction. Antenna gain from directional beams can significantly extend the range (e.g., by 3× when 3.4° beam is used). Moreover, directional beams experience less rate loss across the obstacles, which is consistent with our penetration test. Interestingly, at lowest modulation rate, the range of a 3.4° 60 GHz link (27.5 Mbps) is comparable to WiFi (6 Mbps). This is partly due to the directionality gain, and partly because 802.11ad trades the ultra-wide bandwidth (2.16 GHz) for more robust connectivity at low rate.

Although the preceding experimental setup cannot represent all indoor scenarios, it clearly demonstrates that 60 GHz links can provide reasonable room-level coverage despite NLOS obstructions. Admittedly, compared with 2.4 GHz WiFi, 60 GHz links are more vulnerable to human




Figure 8: Throughput and coverage in indoor environment. Transmit power of 60 GHz links are calibrated to match minimum RSS required for highest data rate at 50 cm. WiFi is calibrated to have 28 dB higher gain than 60 GHz.

blockage, which necessitates new adaptation strategies (Section 4).

3.3 Rate/Range Scaling with Beamwidth

Given the directionality gain, it might be tempting to narrow the beamwidth indefinitely to improve link rate/range, as theoretically range improves quadratically with reduction in beamwidth (contribution via transmitter’s and receiver’s beam). However, ultimately the achievable performance is limited by the FCC’s rule on 60 GHz spectrum, intended to confine the signal radiation for safety and interference management.

Specific to indoor environment, the maximum output power P_t (fed into antenna) is limited to 500 mW (27 dBm). A narrow-beam antenna amplifies the signal and increases its power by G_t dB, i.e., the antenna gain value. But the ultimate $P_t + G_t$, so called *Effective Isotropic Radiated Power (EIRP)*, must be limited to 40 dBm [24]. The joint effect of these two limits is depicted in Figure 9(a). When beamwidth is larger than 11.25°, EIRP increases proportionally as beamwidth decreases, even though P_t is capped to the 500 mW limit. As beamwidth narrows down to below 11.25°, EIRP saturates and is bounded by the 40 dBm limit. Thus, although a further decrease of transmitter beamwidth improves G_t , the P_t has to be decreased to enforce the EIRP limit¹.

However, it should be noted that the EIRP limit only bounds the effective power from the *transmitter*. Link quality can still be boosted as the receiver narrows its beamwidth, thus increasing the receiving antenna gain G_r . To evaluate the ultimate rate/range scaling as a function of beamwidth and under regulation constraints, we follow long-range indoor channel model in [13] to calculate the link budget, considering Friis loss within 1 m distance and log-distance path loss model with exponent 1.6 and a shadowing factor with standard deviation 1.8 dB. Figure 9(b) shows the LOS range scaling under different data rates and beamwidth. In the P_t -bounded region (beamwidth above 11.25°), maximum range increases quadratically as beamwidth decreases. Whereas in the EIRP-bounded region (beamwidth below 11.25°), only the receiver’s narrower beamwidth contributes to RSS increase, and thus the range extension saturates compared with ideal case without EIRP regulation. Note the absolute range values here does not necessarily match all indoor environment. Yet the relative scaling should still hold.

¹Note that the regulation differs in outdoor environment, which allows EIRP to increase by 2 dB for every 1 dB of additional antenna gain, until it reaches a maximum EIRP of 82 dBm [24].




Figure 9: (a) EIRP and maximum output power variation considering FCC’s regulation. (b) Range/rate tradeoffs for different modulation rates and beamwidth.

4. FLEXIBLE BEAMS UNDER LINK DYNAMICS

In this section, we investigate the effectiveness of 802.11ad’s beam searching protocol in establishing high-capacity directional 60 GHz links in various environments, and maintaining the connection under human blockage or device motion. We unveil new challenges in this networking paradigm, and identify potential solution mechanisms that build on unique properties of 60 GHz links.

4.1 Impact of Beam Searching

4.1.1 A primer on 802.11ad beam searching: complexity and overhead

The 802.11ad standard employs a three-phase beam-training mechanism (Figure 10) to search for the optimal transmit and receive beam direction that maximizes link SNR. The *first phase* involves an initiator and a responder performing sector-level sweeping (SLS), *i.e.*, both the transmitter and receiver configure themselves to a *quasi-omni* mode and sweep through discrete directions. The receiver measures the SNR of each direction combination and reports it to the transmitter. This process yields initial coarse-grained sector-shaped beams with low directionality. In the *second phase*, called MID (Multiple sector ID Detection), a group of Tx quasi-omni directions are tested against a number of finer-grained Rx beams within the best Rx sector. This process is used to refine the Tx sector found in the SLS phase, which might have been suboptimal due to the use of quasi-omni Rx directions. A best Tx sector is estimated in this phase which is subsequently used for further refinement in the *third phase* called BC (Beamforming Combining). Here, a set of up to γ fine-grained directions ($\gamma \leq 7$) in the Tx and Rx sector, identified as top candidates in the previous phases, perform pairwise SNR testing.

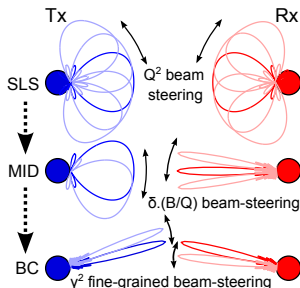


Figure 10: Beamforming training procedure in 802.11ad.

Denote B as the number of fine-grained transmit/receive beam directions, Q as the number of quasi-omni beam directions ($Q = c \cdot B$, where $c < 1$ is a constant defined in the standard [2]), and $\delta (\geq 2)$ as the number of candidate Tx sectors in the MID phase. Then the *timing complexity of the beam training* equals $BF_\tau = Q^2 + \delta \cdot \frac{B}{Q} + \gamma^2$. The SNR testing of each possible beam combination requires sending a separate training frame or sequence, which costs extra channel time. Figure 11 plots the resulting latency for different beamwidths², assuming the same number of antenna elements on the transmitter and receiver.

Due to the exhaustive search in the first phase, the BF_τ grows quadratically ($O(B^2)$) with the number of beam directions. The latency is several orders of magnitude higher than a typical 1KB packet sent at Gbps bit-rate. In Section 4.2, we will evaluate how such latency translates into throughput degradation, depending on how frequently the training process is invoked. By default, 802.11ad triggers beam searching whenever *link outage* occurs (*i.e.*, SNR cannot support the lowest Data PHY rate of 385 Mbps), which could be caused by either human blockage or device motion (beam misalignment).

4.1.2 Link asymmetry: a pathology from flexible beams

Due to the use of discrete phase-shifts on antenna elements, a practical 802.11ad phased-array cannot generate homogeneous beams across all directions [6]. Figure 12 shows an example. Both the AP and STA have 16 antenna elements, with narrowest beamwidth of 11.25° and widest 45° . Due to relative orientation, link SNR is maximized if a 11.25° beam from the STA aligns with a 45° beam from the AP.

Theoretically, even with asymmetric beam patterns, both uplink and downlink should have the same total antenna gains, so that the channel should still be reciprocal. However, we found this no longer holds in the presence of human blockage and environment dynamics.

Specifically, we set up a LOS link with AP orientation of 90° and beamwidth of 90° , and STA orientation of 30° and beamwidth of 22.5° . A human walks and cuts the link across at approximately 2 s intervals while the downlink and uplink packet transmission is ongoing alternately. Beam searching is disabled to isolate the impact of training overhead. The experiment lasts for about 5 minutes and throughput CDF across 50 ms. windows is plotted in Figure 13.

We found that the downlink provides 210 Mbps higher median throughput than uplink. Our further examination reveals that, a narrow receiver beam has higher channel coherence time³ than a wide one *irrespective of the transmitter beamwidth*, resulting in less SNR variation and throughput turbulence in the presence of human movement/blockage. This explains the higher downlink throughput with a narrower beam STA. Such link asymmetry breaks the assumption in many conventional MAC operations (*e.g.*, reciprocity based link quality estimation).

4.1.3 AP discovery using quasi-omni beams

²The relationship between minimum beamwidth Θ and number of steering directions B is given by, $B = (\frac{360}{\Theta})^2$.

³Channel coherence time is the duration over which the channel states remain highly correlated (correlation coefficient ≥ 0.9) [25].

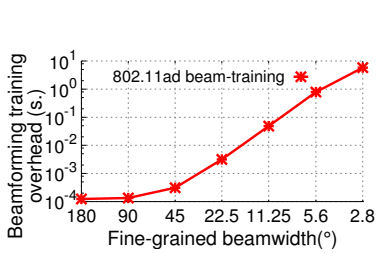


Figure 11: Beamforming training overhead in 802.11ad.




Figure 12: Simulated asymmetric association between AP and STA. Both are using 16-element linear phased-array.




Figure 13: Throughput distribution in presence of human movement/blockage across the link.




Figure 14: AP discovery challenges: (a) Distribution of AP discovery latency using quasi-omni directional antenna. (b) Worst-case discovery latency.

An 802.11ad AP broadcasts beacons periodically, which can be leveraged by newly joining clients for AP discovery and association. To improve the reliability of beaconing, the AP needs to repeat beacons through all its quasi-omni directions. Conceptually, this allows the clients to quickly hear the beacons irrespective of their own beam orientation.

We verify the efficacy of such an AP discovery scheme for both static and mobile clients. Due to the limitation of our experimental platform that requires mechanical steering of horn antenna, we emulate the mobile client using channel traces. Specifically, we move the WiMi receiver and sample spatial channel at discrete points along a pre-defined trajectory, separated 15 cm apart. Each sample requires the client to measure the angular-distribution of RSS by rotating at a granularity of 0.04° . Meanwhile, the AP is fixed at omnimode, with appropriate power control to match the quasi-omni gain. The resulting traces can thus determine the RSS when the client and AP are configured to any beamwidth and direction. A beacon can be received if the RSS exceeds the receiver sensitivity (-78 dBm) corresponding to the lowest modulation level. Following the 802.11ad specification [2], we configure the AP’s quasi-omni beamwidth to ensure its antenna gain is around 15 dB lower than the client’s beam.

Figure 14(a) plots the resulting CDF of AP discovery latency among 6 static locations (10^5 trials each), and along the trajectory of a mobile node. The beacon period equals the default 100 ms. We observe that, unlike legacy WiFi, discovery latency in 802.11ad depends heavily on *location and mobility*. For a static client, the latency ranges from 5 ms to 1.8 s, whereas for a mobile client, it can rise to 12.9 s!

Such long latency might seem counter-intuitive, given that the AP is emulating omni-directional transmission by sending beacons through all quasi-omni directions. To identify the root cause, we use WiMi’s mechanical rotator to test the receiver’s angular RSS distribution while the transmitter is using an omni-directional antenna. It is commonly assumed that, no matter which direction it points to, the receiver can




Figure 15: Examples of omni-directional transmission’s AOA pattern captured at two different receiver’s location.

hear an omni-directional transmitter. Besides AP discovery, such assumption is often leveraged in directional MAC protocols as a means of broadcast signaling [26]. However, example tests in two locations (Figure 15) disproves the perception. The receiver’s angular RSS is sparsely distributed and largely pronounced on several narrow directions spanning a few degrees. This is likely attributed to the presence of a few strong reflecting objects indoor, and thus independent of the beamwidth of transmitter. Hence, unless both the transmitter’s and receiver’s quasi-omni beams are tuned to cover the reflecting directions, the beacon will be missed, resulting in prolonged discovery time.

We proceed to evaluate the scenarios with different number of antenna elements (*i.e.*, beamwidths). For each location, the measured angular RSS distribution is sectorized and calibrated in accordance to the quasi-omni beamwidth generated by a given number of antenna elements. In each experiment, the STA picks a random sector to listen to, and moves to the next sector if the AP beacon is not heard within the 100 ms beacon interval. The experiment is repeated over 10^5 random STA orientations. From the results (Figure 14(b)), we see that the worst-case latency grows proportionally with the number of antenna elements. Mobile nodes experience significantly higher latency than static ones (*e.g.* 5 s and 18 s, for 64 and 128 antenna elements respectively), which will clearly hinder the adoption of 60 GHz mobile devices.

The deficiency of quasi-omni transmission/reception can affect a wide range of signaling mechanisms in 802.11ad. For example, to resolve inter-cell interference in TDMA mode, an AP needs to send a *Quieting Adjacent BSS (QAB)* message to adjacent APs to suppress their transmission, in a similar way to periodic beaconing. Latency of such QAB signaling would be similar to the beaconing, thus resulting in substantial interference before the APs can respect each other. Although enforcing omni-directional signaling at both ends may alleviate the situation, it will substan-




Figure 16: Performance of beam-steering algorithms in presence of user motion. (a) Packet throughput. (b) Steering latency. Measurements are averaged over 1 s window.

tially reduce range (Section 3.2). Alternative PHY-layer technologies, *e.g.* spread-spectrum, that trades bit-rate for range, may be a complementary solution, but is beyond the scope of this work.

4.2 Efficient Beam Searching based on AOA Sparsity

Although the sparse clustering of angular RSS plagues AP discovery, we found it can be exploited to curtail the overhead in beam searching. Below we propose a heuristic algorithm that can find near optimal beam direction without exhausting the full beam space.

Suppose two stations STA_a and STA_b search for the beam configuration that maximizes link SNR. Our basic idea is to allow STA_a to transmit beam training frames in *omni-directional* mode, while STA_b sweeps its receive beam directions using fine-beams (instead of quasi-omni) and evaluates the SNR. In this way, a consistent sparse AOA pattern is maintained throughout the training procedure, and the receiver is guaranteed to rendezvous with the angle with strongest signals. The role gets reversed when STA_a wants to identify its own beam direction. Overall, this simple algorithm can reduce search complexity to $O(B)$, substantially lower than 802.11ad (Section 4.1).

To validate the mechanism, we set up 10 different links each using 3 different beamwidths and collect RSS traces for 15 min. Multiple people move around the room and unintentionally block the links. From the trace data of RSS changes we evaluate the frequency of invoking an 802.11ad beam-searching. We then run the MAC emulator in WiMi that accounts for both the link rate and overhead. The results in Figure 16(a) show that, for 802.11ad, narrower beamwidth actually provides worse performance as antenna gain (SNR improvement) is nullified by the beam-steering overhead, even with an intermediate beamwidth of 22.5°. Figure 16(b) further contrasts the latency of beam-steering. Latency of the sparse-AOA based algorithm shows a linear growth while the 802.11ad has a quadratic growth as beamwidth narrows down (thus increasing beam search space), which explains its relatively higher throughput for narrower beams. It should be noted that the maximum link distance during beam searching may differ from that in data transmission. Depending on the sum of transmit/receive antenna gains during the quasi-omni mode SLS, 802.11ad’s beam searching range can be either higher or lower than sparse-AOA.

4.3 Beam Switching to Overcome Blockage

Besides aligning the transmitter and receiver’s beam directions to maximize link SNR, a crucial functionality of




Figure 17: Link improvement due to beam dilation *w.r.t.* normalized human blockage position.

beam searching is to overcome human blockage and track mobile users, which we evaluate in this section.

4.3.1 Adapting Beamwidths to Overcome Blockage

One intuitive solution to human blockage is dilating the transmitter/receiver’s beamwidth, such that part of the beam circumvents the human body. We now test the reliability of this approach in a controlled setting, with one LOS link (7 m distance). The original beamwidth is 3.4° and when the link is blocked, the transmitter switches to 19.2° (corresponding to its widest quasi-omni beam) without changing its direction. Figure 17 plots the link rate when the human blocks different position of the link (normalized *w.r.t.* link distance). Evidently, when the original link was operating at high rate of 4.6 Gbps (with a minimum SNR of 27 dB), the link can still sustain connectivity, but mostly using the lowest 27.5 Mbps Control PHY rate, and occasionally upgraded to a higher rate when blockage occurs in mid-link. In contrast, when the original SNR was low (12 dB with minimum Data PHY rate of 385 Mbps), no link connectivity is possible irrespective of the blockage position.

Therefore, under human blockage, beam dilation can help maintain link connectivity, but only if the original link SNR is high, and when blockage is sufficiently far away from the transmitter/receiver, such that only part of the beam is blocked. Regardless, the link rate will degrade substantially.

4.3.2 Steering Beam Direction

Alternatively, the transmitter can steer its beam direction to find a detour path that bounces off reflective objects. IEEE 802.11ad invokes this scheme (Section 4.1) automatically whenever the link quality suffers. To evaluate its efficacy, we define a *beam-steering effectiveness* metric η as follows,

$$\eta = (R_1 - R_0)/R_0 \quad (2)$$

where R_0 denotes the LOS link rate before blockage and R_1 means the best achievable rate after beam-steering. η should thus satisfy, $-1 \leq \eta < \infty$. $\eta = -1$ if none of the secondary paths can establish connectivity. A positive η implies that secondary path after beam-steering provides a higher rate than the original LOS path. This can happen when the LOS path itself is weaker. Note that, RSS increase of as small as 1 dB can improve η by 0.3 on average⁴. Here we only focus on the rate change before/after beam steering to isolate the impact of searching overhead.

Figure 18(a) and 18(b) plot the η in an office and corridor environment, respectively. For each beamwidth setting, we test 10 links with transmitter and receiver randomly placed inside the test region, link distance ranging from 3 m to 7 m.

⁴Calculated using the rate table 2 in Appendix C and considering an operating RSS in the range -78 to -53 dBm.

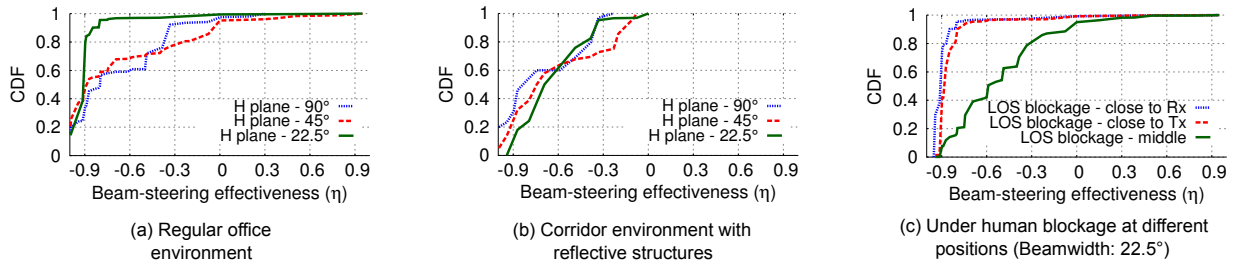


Figure 18: Beam-steering effectiveness (η) in indoor environment.

We see that the η of a narrow beam (22.5°) is remarkably lower than that of wider beams (45° & 90°). Recall the partitions and drywalls in the office environment mostly pass instead of reflecting 60 GHz signals (Section 3.2), whereas candidate reflecting surfaces (mostly from 4 tall metal cabinets and book shelves) are sparsely distributed. Thus, the narrow beam can rarely find an opportunistic path to detour towards the receiver. Amid its lower rate, a wider beam is more resilient, as long as it can be steered to the best secondary path.

Reflective paths in the corridor environment show up differently — all beamwidths exhibit a reasonable level of η , owing to a more open space and existence of brick walls and concrete poles. For 45° and 90° , around 6% and 20% of links experience zero rate even after beam-steering, primarily because of the reflection loss and low antenna gain. Note that in both environments, we are only steering the beams around the azimuthal plane. Depending on reflectivity of ground/ceiling, the effects on the elevation dimension would show up following the same principles.

Figure 18(c) further plots the impact of human blocking position on η , for a 22.5° beam in corridor. We see that η can be much higher if blockage occurs in mid-link, in which case the beam is more likely to find a detour path. When blockage is close to either transmitter or receiver, η is degraded to around -0.9 , *i.e.*, 90% of rate is lost.

In summary, *despite higher rate, a narrow beam is more susceptible to blockage as it owns fewer escape paths. Beam-steering can be effective in reflective environment, but barely helps when blockage occurs in close proximity of transmitter/receiver, e.g., human holding a device.*

4.4 Assisting Link Recovery via mmWave Sensing

The above micro-benchmark measurement directly compares possible beam dilation/steering options to overcome blockage. A practical protocol has to choose beam pattern based on its observation of link status. The 802.11ad adopts an *SNR maximization* strategy that tracks the beam pattern to maximize link SNR, which we evaluate below.

For the device motion experiment, we program WiMi’s motion controller and create random orientation/position trajectory to emulate typical scenarios when user is holding a mmWave mobile device within the AP’s LOS. Real-time RSS traces are collected for a 3.4° link and corresponding 19.2° quasi-omni mode used for beam dilation. The throughput distribution across 1800 time windows (each 50 ms) is plotted in Figure 19(a). Due to dramatic link quality variation, the 802.11ad beam searching is frequently invoked, resulting in substantial overhead. Consequently, dilating the




Figure 19: Throughput of beam-steering and beam-dilation approach: (a) Device motion. (b) Human blockage.

beam to quasi-omni leads to $3\times$ higher throughput, despite its lower bit-rate (when steering overhead is not considered).

For the human blockage experiment, the transmitter and receiver are static and within LOS. A human walks and cuts the link across repeatedly for 90 seconds. Due to intermittent blocking effects, a dilated beam degrades throughput to the lowest value for more than 50% of the times (note this differs from the experiment in Section 4.3.1 where blockage is permanent). In contrast, median throughput of the narrower 3.4° beam is an order of magnitude higher. Beam steering is invoked infrequently — only when the beam is being cut, followed by periodic update in the beginning of beacon interval (Section 2.1). Thus, beam steering overhead is much smaller compared with the device motion case, and SNR-maximizing beam-steering can still be effective.

To summarize, *naïve SNR maximization leads to substantial overhead and low throughput for mobile devices, but can still be effective for a static devices even under human blockage. Thus, it is necessary to execute this metric according to the root cause of link dynamics.*

The natural question thus follows is: how can we diagnose the link outage and invoke proper protocol reaction? We observe that the inherent sensitivity of mmWave links can be harnessed to realize the diagnosis.

Figure 20(a) shows an example of RSS variation over short period of time due to human blockage (cutting the link) and device motion of a 3.4° link. Clearly, human blockage causes a more regular pattern of RSS change, which can serve as a signature to distinguish it from device motion. Let $RSS(t)$ denote the receiver’s RSS at time t . We capitalize on the second order statistics of RSS values at the receiver side gathered over a time window of T , calculated as follows,

$$\partial^2 RSS(t, T) = \frac{\partial^2}{\partial t^2} [RSS(t), \dots, RSS(t + T)], \quad (3)$$

$$N_{\partial^2 RSS} = \text{Var}\{\partial^2 RSS(t, T)\}$$

where $\text{Var}\{\cdot\}$ denotes the variance of sequence of observations. A window is classified as blockage if $N_{\partial^2 RSS}$ falls below




Figure 20: (a) Example of RSS variation due to human blockage and device motion. (b) Accuracy of classification for human blockage and device motion. Beamwidth: 3.4° .

a threshold Θ_{block} , and motion if $N_{\partial^2 RSS}$ exceeds an upper threshold Θ_{motion} . We then accumulate the confidence of classification over time until it rises above 0.95, or exceeds a detection latency bound t_{pred} (see Algorithm 1). The two thresholds are empirically set to 0.002 and 0.2. Since they are far away from each other, we found the detection performance is insensitive to minor parameter adjustment.

To verify this approach in realistic environment, we repeat the human blockage and device motion each for 50 times. We evaluate the classification accuracy for a given t_{pred} , computed over all 50 trials. From the Figure 20(b), we see that the accuracy increases sharply beyond a certain period of confidence accumulation. Given 30 ms or above, both motion and blockage can be classified with close to 100% accuracy. Although the 30 ms latency may leave the link in a temporarily outage/low-throughput state, it helps making a throughput maximizing choice immediately afterwards, instead of relentlessly seeking for SNR maximization, which may result in multiple folds of throughput reduction (Figure 19).

Algorithm 1 Detecting device motion and human blockage

- 1: Initialize $t = 0$, $T = T_{win}$, $C = 1$, $\alpha_{\{block, motion\}} = 0$.
 - 2: **while** $\alpha \leq 0.95$ /* 95% confidence */
 - 3: Calculate the variance of second order statistics of RSS samples $N_{\partial^2 RSS}$, from Equation 3.
 - 4: **if** $N_{\partial^2 RSS} \leq \Theta_{block}$ **then** $\alpha_{block} = \alpha_{block} + 1$;
 - 5: **else if** $N_{\partial^2 RSS} \geq \Theta_{motion}$ **then** $\alpha_{motion} = \alpha_{motion} + 1$;
 - 6: **end if**
 - 7: **if** $\alpha_{block} > \alpha_{motion}$ **then** $\alpha = \alpha_{block}/C$;
 - 8: **else** $\alpha = \alpha_{motion}/C$;
 - 9: **end if**
 - 10: $T = T + T_{win}$; $C = C + 1$;
 - 11: **if** $T \geq t_{pred}$ /* Is it too late for prediction? */
 - 12: Update t , $T = T_{win}$, $C = 1$, $\alpha_{\{block, motion\}} = 0$;
 - 13: **end if**
 - 14: **end while**
-

Note that, the above device motion experiments were conducted in LOS. In practical usage scenario of mobile mmWave devices, human blockage and device motion can occur together (*e.g.* human holding a device and blocking the AP), where neither beam-steering nor dilation is guaranteed to work. 802.11ad has proposed a fast session transfer mechanism to relocate the link to 2.4/5 GHz, thus salvaging the connection. Measurement and evaluation of such mechanisms is left as our future exercise.




Figure 21: Distribution of spatial reuse factor (β) with 3 different beamwidths.

5. SPATIAL REUSE BETWEEN FLEXIBLE-BEAM LINKS

We now study the spatial reuse between narrow-beam 60 GHz links in both static and dynamic scenarios.

5.1 Imperfect Spatial Reuse Between Static Links

Existing modeling/measurement of spatial reuse in flexible-beam directional antenna networks presented bifurcated views: measurement of switched-beam WiFi radios in indoor environment [27] showed that directional beams provide little advantage of spatial reuse due to strong multipath reflections; on the other hand, ray-tracing modeling of outdoor 60 GHz networks [28] claims that the narrow directional links can be abstracted as pseudo-wire with no leakage interference. We now measure the realistic spatial reuse in indoor environment, characterized by a spatial reuse factor, defined as [27]:

$$\beta = \frac{\text{Sum rate of concurrent links}}{\text{Average rate of isolated links}} \quad (4)$$

Ideally, two coexisting links can achieve $\beta = 2$, should there be no mutual interference. The smaller β is, the stronger mutual interference becomes, and the worse spatial reuse.

Distribution of spatial reuse factor. We evaluate β in the office environment. For each beamwidth configuration, we randomly create 10 pairs of links, measure the SINR of each link, and then translate it into link rate (using Table 2) and β . Node locations are created under an amenable configuration such that interfering transmitter does not point directly to unintended receiver. Figure 21 plots the CDF of β across all pairs. In general, wider beams experience smaller spatial reuse. For a 3.4° narrow beam, perfect spatial reuse is achieved in 75% of cases. Yet there still exist a sizable fraction of cases with $\beta < 2$. Therefore, unlike 2.4 GHz networks [27], perfect spatial reuse is practical in the common cases for highly directional 60 GHz beams indoor, but leakage or reflected interference is still non-negligible, and thus the pseudo-wire abstraction does not hold.

Impact of side lobes. Existing modeling and simulation typically use a fan-shape to represent the spatial footprint of a 60 GHz link [28]. However, practical 60 GHz phased-array antennas inevitably have spurious side lobes that leak signals [6]. To evaluate the impact on spatial reuse, we place a static receiver on 7 different positions (out of the 10 locations in the previous experiment) and an interferer on 3 locations.

We measure interference on the unintended receiver locations, with and without side lobe effect. The transmitter is a COTS device [15] that uses a main beamwidth of 30° and has side lobes as shown in Figure 22(b). For contrast purpose, we also emulate the scenario without side lobes by enclosing the transmitter's antenna with two 50 dB mmWave




Figure 22: (a) Measured spatial reuse factor (β) using 30° beamwidth across 7 different locations. (b) Example of measured beam pattern from the Wilocity 60 GHz radio.

absorber boards, which form a 30° conical-shape opening and attenuate the side lobes to a negligible level.

Figure 22(a) plots the resulting β factor across 7 receiver locations. We see that side lobes can degrade β by 6% to 25%, translating into 250 Mbps to 700 Mbps of aggregate link throughput reduction. Thus, *to be precise, an interference model must take into account such imperfectness of practical phased-array antennas, instead of abstracting links as fan shapes or pseudo-wires.*

5.2 Spatial Reuse Under Link Dynamics

To maximize spatial reuse, an 802.11ad AP runs an interference aware scheduler that builds a conflict graph between directional links. While one link is transmitting, others can sense and periodically report the overheard RSS (interference) level to the AP. The AP then groups those links with close-to-zero mutual interference into the same TDMA slot to maximize spatial reuse, and regroup if interference recurs. One natural question is: will such interference-aware scheduling be able to cope with link dynamics?

We answer the question through a micro-benchmark experiment illustrated in Figure 23(a). The setup contains two links Tx1 \rightarrow Rx1 and Tx2 \rightarrow Rx2, each with 22.5° beamwidth. Initially both are static and orthogonal to each other, and thus grouped for concurrent transmission. Then, we move Rx1 at walking speed, passing by Rx2. While Tx1 is tracking Rx1, its beam covers and interferes with Rx2.

An 802.11ad AP handles such situations by recomputing a schedule in every beacon period (100 ms by default). We compare this scheme with two alternatives: (i) a fine-grained scheduler that requires STAs to report to the AP immediately whenever it wants to execute a beam-steering procedure (with another STA) so that the AP can take appropriate action if the conflict graph is changed due to new beam direction; (ii) a TDMA scheduler that intentionally allocates a separate slot to the mobile link. We plot the resulting CDF of throughput across every 50 ms window in Figure 23(b). 802.11ad’s beacon-level granularity is too low to respond to node mobility, and thus its throughput is much lower than the isolated TDMA in majority of cases. On the other hand, the fine-grained scheduler incurs too much feedback overhead, and thus even lower throughput. The isolated TDMA performs best on average, yet its throughput in a small fraction of high-end cases is lower, mainly because it conservatively separates the two links even before they become conflicting.

To summarize, *an interference-aware scheduler must balance responsiveness and overhead. To make the best tradeoff, a mobile link should be opportunistically isolated from others and regrouped when appropriate.* The motion detection

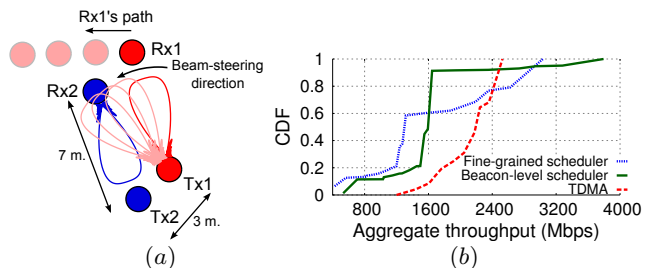


Figure 23: (a) Example of interference map change in presence of mobile client. (b) Distribution of aggregate throughput under the influence of mobile client.

algorithm (Section 4.4), combined with interference prediction (leveraging movement direction and relative position of links), hold potential to make such cognitive decisions.

6. RELATED WORK

60 GHz measurement and modeling. Prior literature on 60 GHz measurement has been extensively focused on characterizing and modeling the channel propagation profile, both indoor [10, 11, 13] & outdoor [29, 30]. Dedicated channel sounder hardware is used to extract statistical properties, such as path loss exponent, delay spread and AOA distribution that can be fit into a general analytical model. Few works measured [31] or simulated [32] the impacts of human blockage, again emphasizing RSS attenuation of a fixed beam, instead of behaviors of flexible-beam protocols.

Recently, Zhu *et al.* [9] used COTS 60 GHz radios for a reality-check of pico-cellular networks, focusing on outdoor coverage and transport-layer throughput. Appealing prospects were revealed: human can rarely block links owing to base-station height (above 6 m); beams can bounce off concrete buildings to reach users; interference can be avoided through base-station coordination. Our measurement found that many of the conclusions no longer hold in indoor 802.11ad WLANs. In particular, human blockage and device mobility pose grand challenges, and the efficacy of beam steering becomes highly sensitive to environment. More importantly, we identify protocol-level issues in operating flexible beams, such as link asymmetry, association latency, beam searching overhead, *etc.* Tie *et al.* [8] also conducted indoor 60 GHz measurement. But similar to [9], only coarse-grained TCP level statistics can be collected due to lack of access to COTS devices’ MAC/PHY layers.

Directional-antenna networking over lower frequencies. Research issues in mmWave networking share certain similarities with the vast literature of directional antenna networking, which focused on MAC/routing protocols for ad-hoc networks with directional antennas (see [26] and the references therein). Steerable phased-array antennas have also been employed to improve spatial reuse between indoor WiFi links [12], enable high-capacity communication and robust outdoor vehicular networking [33]. The 60 GHz 802.11ad and 802.15.3c standards partly assimilated these research ideas, particularly in its hybrid MAC design. However, mmWave networking faces many unique challenges imposed by ultra-narrow beams (an order of magnitude narrower than those assumed in microwave-band), sensitivity to path loss/blockage, *etc.*, which have been testified in our contrasting experiments with WiFi (Section 3.2). In addi-

tion, unique capabilities of 60 GHz devices, *e.g.*, fine-grained beam switching, blockage/motion sensing, smaller spatial footprint, deserve new system design and measurement investigation.

7. CONCLUSION

In this paper, we have presented a link-level measurement of indoor 60 GHz networks using a software-radio platform. The measurement reveals new challenges pertinent to MAC protocols involving ultra-directional, flexible 60 GHz beams. In particular, 60 GHz links become highly sensitive to human blockage and device motion, posing non-trivial trade-offs between link quality and responsiveness, especially *w.r.t.* the beam searching and interference-aware scheduling. On the other hand, abundant opportunities exist. For example, a 60 GHz ultra-directional link can cover beyond a single room (unlike the common perception of pseudo-optical propagation). Its sensitivity can be harnessed to diagnose link outage, thereby facilitating judicious protocol reactions. We believe a new class of protocols need to be designed to incorporate such challenges/opportunities.

Acknowledgement

We sincerely thank the anonymous reviewers for their valuable comments and feedback. The work reported in this paper was supported in part by the NSF under Grant CNS-1318292, CNS-1343363, CNS-1350039 and CNS-1404613.

8. REFERENCES

- [1] Qualcomm Inc., “The 1000x data challenge,” <https://www.qualcomm.com/1000x>, 2014.
- [2] IEEE Standards Association, “IEEE Standards 802.11ad-2012: Enhancements for Very High Throughput in the 60 GHz Band,” 2012.
- [3] —, “IEEE Standards 802.15.3c-2009: Millimeter-wave-based Alternate Physical Layer Extension,” 2009.
- [4] X. Zhou, Z. Zhang, Y. Zhu, Y. Li, S. Kumar, A. Vahdat, B. Y. Zhao, and H. Zheng, “Mirror Mirror on the Ceiling: Flexible Wireless Links for Data Centers,” in *Proc. of ACM SIGCOMM*, 2012.
- [5] T. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. Wong, J. Schulz, M. Samimi, and F. Gutierrez, “Millimeter Wave Mobile Communications for 5G Cellular: It Will Work!” *IEEE Access*, vol. 1, 2013.
- [6] T. S. Rappaport, R. W. H. Jr., R. C. Daniels, and J. N. Murdock, *Millimeter Wave Wireless Communications*. Prentice Hall, 2014.
- [7] H. Pan, M. Park, and H. Alavi, “Array Analysis Using High Efficiency mm-wave Antenna for Gigabit Plus Throughput Wireless Communication,” in *IEEE Antennas and Propagation Society International Symposium*, 2009.
- [8] X. Tie, K. Ramachandran, and R. Mahindra, “On 60 GHz Wireless Link Performance in Indoor Environments,” in *PAM*, 2011.
- [9] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson, U. Madhow, B. Y. Zhao, and H. Zheng, “Demystifying 60GHz Outdoor Picocells,” in *Proc. of ACM MobiCom*, 2014.
- [10] H. Xu, V. Kukshya, and T. Rappaport, “Spatial and Temporal Characteristics of 60-GHz Indoor Channels,” *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 3, 2002.
- [11] C. Anderson and T. Rappaport, “In-Building Wideband Partition Loss Measurements at 2.5 and 60 GHz,” *IEEE Transactions on Wireless Communications*, vol. 3, no. 3, 2004.
- [12] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste, “DIRC: Increasing Indoor Wireless Capacity Using Directional Antennas,” in *Proc. of ACM SIGCOMM*, 2009.
- [13] P. F. M. Smulders, “Statistical Characterization of 60-GHz Indoor Radio Channels,” *IEEE Transactions on Antennas and Propagation*, vol. 57, no. 10, 2009.
- [14] X. Zhang and P. Ramanathan, “WiMi Project Website,” 2014. [Online]. Available: <http://xyzhang.ece.wisc.edu/wimi>
- [15] “Wilocity 802.11ad Multi-Gigabit Wireless Chipset,” <http://wilocity.com>, 2013.
- [16] Rice University, “Wireless Open-Access Research Platform,” <http://warp.rice.edu/trac/wiki>, 2013.
- [17] Ettus Inc., “USRP,” <http://www.ettus.com>.
- [18] K. Tan, J. Zhang, J. Fang, H. Liu, Y. Ye, S. Wang, Y. Zhang, H. Wu, W. Wang, and G. M. Voelker, “Sora: High performance software radio using general purpose multi-core processors,” in *Proc. of USENIX NSDI*, 2009.
- [19] “60 GHz Transmit/Receive Development Systems,” <http://www.pasternack.com>, 2014.
- [20] TDK RF Solutions Inc., “Absorber for Microwave and Millimeter Wave Test Chambers,” 2014.
- [21] Y. Zhu, X. Zhou, Z. Zhang, L. Zhou, A. Vahdat, B. Y. Zhao, and H. Zheng, “Cutting the Cord: a Robust Wireless Facilities Network for Data Centers,” in *Proc. of ACM MobiCom*, 2014.
- [22] IEEE 802.11 NG60 Study Group, “Next Generation 60 GHz (NG60),” 2012.
- [23] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.
- [24] “FCC 13-112,” <https://apps.fcc.gov/edoc-public/attachmatch/FCC-13-112A1.pdf>, 2013.
- [25] X. Xie and X. Zhang, “Adaptive Feedback Compression for MIMO Networks,” in *Proc. of ACM MobiCom*, 2013.
- [26] O. Bazan and M. Jaseemuddin, “A Survey On MAC Protocols for Wireless Adhoc Networks with Beamforming Antennas,” *IEEE Communications Surveys and Tutorials*, vol. 14, no. 2, 2012.
- [27] S. Lakshmanan, K. Sundaresan, S. Rangarajan, and R. Sivakumar, “The Myth of Spatial Reuse with Directional Antennas in Indoor Wireless Networks,” in *PAM*, 2010.
- [28] S. Singh, R. Mudumbai, and U. Madhow, “Interference Analysis for Highly Directional 60 GHz Mesh Networks: the Case for Rethinking Medium Access Control,” *IEEE/ACM Transactions on Networking*, vol. 19, no. 5, 2011.
- [29] E. J. Violette, R. H. Espeland, and G. R. Hand, “Millimeter-wave Urban and Suburban Propagation Measurements Using Narrow and Wide Bandwidth Channel Probes,” *NASA Tech Report*, 1985.
- [30] T. Rappaport, F. Gutierrez, E. Ben-Dor, J. Murdock, Y. Qiao, and J. Tamir, “Broadband Millimeter-Wave Propagation Measurements and Models Using Adaptive-Beam Antennas for Outdoor Urban Cellular Communications,” *IEEE Transactions on Antennas and Propagation*, vol. 61, no. 4, 2013.
- [31] S. Collonge, G. Zaharia, and G. Zein, “Influence of the Human Activity on Wide-Band Characteristics of the 60 GHz Indoor Radio Channel,” *IEEE Trans. on Wireless Comm.*, vol. 3, no. 6, 2004.
- [32] M. Park and H. K. Pan, “Effect of Device Mobility and Phased Array Antennas on 60 GHz Wireless Networks,” in *ACM mmCom*, 2010.
- [33] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das, “MobiSteer: Using Steerable Beam Directional Antenna for Vehicular Network Access,” in *Proc. of ACM MobiSys*, 2007.




Figure 24: WiMi waveguide antenna beam pattern: (a) Horizontal plane. (b) Vertical plane.




Figure 25: Examples of simulated beam pattern vs. beam pattern generated using mmWave absorbers.

APPENDIX

A. EMULATING FLEXIBLE BEAM PATTERNS ON WIMI

We exploit WiMi’s quasi-omni directional antenna waveguide, together with multiple mmWave absorbers [20], to create a variety of beam patterns. The RF front-end is pre-built with a waveguide that provides $\sim 180^\circ$ beamwidth, which we have verified by placing a WiMi receiver around and sniffing its angular RSS distribution (Figure 24). The absorbers are placed in front of the waveguide in conical shape to create the desired beam patterns. Examples of beam pattern generated using this emulation approach is shown in Figure 25, which closely matches the ideal beam pattern. Thus, the emulated beams suffice for evaluating the impact of different beamwidths. We have attached the absorbers to the RF front-end, and mount them together on WiMi’s motion control system, which allows for programmable pan/tilt/linear movement, thus ready for emulating beam steering and device motion.

B. EMULATING 60 GHZ MIMO ON WIMI

Due to lack of a commercial 60 GHz MIMO array, we create a virtual array of antennas to emulate a 2×2 MIMO node on WiMi.

A WiMi radio has only one Tx/Rx RF chain, supporting one Tx/Rx antenna. We overcome this constraint by creating a virtual MIMO Tx/Rx through antenna displacements. Specifically, we use the motion controller to slide an omni-directional Tx antenna to 2 different positions, separated by half-wavelength (2.5 mm). Meanwhile, the Rx antenna measures the channel gain (phase and magnitude). The measurement is then repeated for 2 Rx antenna positions, to obtain a 2×2 matrix representing the MIMO channel.

To ensure the matrix represents a “snapshot” of the MIMO channel, all the 4 element-wise channels must remain stable across the entire measurement. To validate this condition is satisfied, we first measure the channel stability, characterized by the *coherence time*. Correlation between two instances of channel samples, separated by T , is defined as [23]:

$$\mathcal{K}(T) = \frac{L}{L-T} \cdot \frac{|\sum_{t=0}^{L-T-1} h(t)h^*(t+T)|}{\sum_{t=0}^{L-1} |h(t)||h^*(t)|} \quad (5)$$

where L is the total length of the sequence of sampled channel gains. Coherence time is represented by an instance $T_{0.9}$ when correlation drops below 0.9, *i.e.*, $T_{0.9} = \mathcal{K}^{-1}(0.9)$. Figure 26 shows that the channel remains highly stable within 50 seconds window, whereas our MIMO channel matrix measurement only takes around 35 to 40 seconds (due to mechanical adjustment), which verifies the effectiveness of our MIMO emulation.

So, how can we measure MIMO capacity based on the emulator? Lets assume H_f forms the channel matrix between the N_T transmit and N_R receive antennas for frequency subcarrier f . Further if $\lambda_1 \geq \lambda_2 \geq \dots \lambda_{\text{rank}(H_f)}$ are the ordered eigen values of the channel matrix H_f . Then, the capacity of the MIMO channel is given by [23],

$$\mathbf{C} = \mathbb{E}_H \left[\sum_{i=1}^{\text{rank}(H_f)} \log_2 \left(1 + \frac{P}{N_0 N_T} \cdot \lambda_i^2 \right) \right] \quad (6)$$

where P denotes the combined output transmit power from N_T transmit antennas and is distributed equally among them. N_0 specifies a constant noise floor across all the received antennas.

In our trace driven emulation, we placed the transmitting and receiving antenna in LOS condition and collected channel matrix traces between the N_T and N_R antenna elements by using virtual sets of omni-directional antennas as described earlier. The channel matrix is collected across 256 different frequency subcarriers in order to capture frequency diversity within the environment. We derived the channel capacity of 60 GHz in the above scenario under a fixed power from the transmitting device according to the Equation 6.

C. MODULATION AND CODING SCHEMES IN IEEE 802.11AD

IEEE 802.11ad introduces three heterogeneous PHY layers dedicated to different application scenarios. The *Control* PHY is designed for very low SNR operation prior to beamforming. The *Single Carrier (SC)* PHY enables power efficient and low complexity transceiver implementation by using less robust coding protocol and using up to 16-QAM complex modulation scheme. The *OFDM* PHY provides high performance in frequency selective channels achieving highest data rates. The data rates, corresponding receiver sensitivity and coding schemes for *Control* and *SC* PHY are shown in Table 2. This table is used to map measured RSS into packet bit-rate throughout our experiments, in a similar way to existing studies [4, 9].

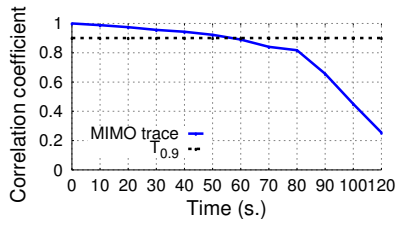


Figure 26: Channel stability during MIMO trace collection.

Modulation	Coding Rate	Receive sensitivity (dBm)	Data Rate (Mbps)
DBPSK	1/64	-78	27.5
$\pi/2$ -BPSK	1/4	-68	385
$\pi/2$ -BPSK	1/2	-66	770
$\pi/2$ -BPSK	5/8	-65	962.5
$\pi/2$ -BPSK	3/4	-64	1155
$\pi/2$ -BPSK	13/16	-62	1251.25
$\pi/2$ -QPSK	1/2	-63	1540
$\pi/2$ -QPSK	5/8	-62	1925
$\pi/2$ -QPSK	3/4	-61	2310
$\pi/2$ -QPSK	13/16	-59	2502.5
$\pi/2$ -16QAM	1/2	-55	3080
$\pi/2$ -16QAM	5/8	-54	3850
$\pi/2$ -16QAM	3/4	-53	4620

Table 2: 802.11ad Control and Single Carrier rates.