

Steering with Eyes Closed: mm-Wave Beam Steering without In-Band Measurement

Thomas Nitsche^{†‡}, Adriana B. Flores*, Edward W. Knightly*, and Joerg Widmer[†]

[†]IMDEA Networks Institute, Madrid, Spain

*ECE Department, Rice University, Houston, USA

[‡]Univ. Carlos III, Madrid, Spain

Abstract—Millimeter-wave communication achieves multi-Gbps data rates via highly directional beamforming to overcome pathloss and provide the desired SNR. Unfortunately, establishing communication with sufficiently narrow beamwidth to obtain the necessary link budget is a high overhead procedure in which the search space scales with device mobility and the product of the sender-receiver beam resolution. In this paper, we design, implement, and experimentally evaluate *Blind Beam Steering (BBS)* a novel architecture and algorithm that removes in-band overhead for directional mm-Wave link establishment. Our system architecture couples mm-Wave and legacy 2.4/5 GHz bands using out-of-band direction inference to establish (overhead-free) multi-Gbps mm-Wave communication. Further, BBS evaluates direction estimates retrieved from passively overheard 2.4/5 GHz frames to assure highest mm-Wave link quality on unobstructed direct paths. By removing in-band overhead, we leverage mm-Wave’s very high throughput capabilities, beam-width scalability and provide robustness to mobility. We demonstrate that BBS achieves 97.8% accuracy estimating direction between pairing nodes using at least 5 detection band antennas. Further, BBS successfully detects unobstructed direct path conditions with an accuracy of 96.5% and reduces the IEEE 802.11ad beamforming training overhead by 81%.

I. INTRODUCTION

The IEEE 802.11ad millimeter-wave WiFi standard promises data rates of up to 7 Gbps via high gain antenna arrays used to overcome the increased attenuation of 20-40 dB at the mm-Wave band [1], [2], [3], [4]. With high directionality, the sender and receiver must match the potentially narrow directions of their respective beams. IEEE 802.11ad discretizes the search space and divides the antenna radiation sphere into as many as 128 virtual sectors supporting beamwidths less than 3 degrees. A two stage process is used to select the best sector: first, an exhaustive search to select the best transmit sector is performed by both transmitter and receiver. Second, directionality gain is increased by fine-tuning antenna settings at both the transmitter and receiver [4].

The search space and overhead of this process directly scales with the product of the number of nodes and their transmit and receive sectors. While, this is acceptable for mm-Wave use cases like wireless HDMI with mostly static transmitters, the next generation of mm-Wave systems targets the mobile device [3], [5], [6]. Device mobility breaks the beam adjustment and requires constant retraining, significantly increasing the beam forming overhead. This holds especially for device rotation. From experiments with a 7 degree beam width system we

found that a mere misalignment of 18 degree reduces the link budget by around 17 dB. According to IEEE 802.11ad coding sensitivities [3], this drop can reduce the maximum throughput by up to 6 Gbps or break the link entirely. This misalignment is easily reached multiple times per second by a user holding the device, causing substantial beam training overhead to enable multi-Gbps throughput for mm-Wave WiFi.

In this paper, we design, implement, and experimentally evaluate **Blind Beam Steering (BBS)**, a system to steer mm-Wave beams by replacing in-band trial-and-error testing of virtual sector pairs with “blind” out-of-band direction acquisition. BBS exploits diverse frequency band characteristics of multi-band WiFi devices compliant with the IEEE 802.11ad standard. To this aim, it utilizes legacy 2.4/5 GHz bands to estimate the direction for pairing nodes from passively overheard frames, which does not incur any additional protocol overhead. BBS uses out-of-band direction interference mechanisms that list received signal energy over an azimuthal receive spectrum. A history of these direction estimates is maintained for every potential pairing device in the network. Whenever a link is to be beam trained, the history is queried first for a valid direction estimate that can replace in-band training.

However, the direction information derived from the legacy-band does not always correspond to the best mm-Wave band direction due to multipath and noise effects. Further, frequency dependent transmission coefficients may lead to direction estimates that steer the beam into an obstacle, causing high mm-Wave attenuation. BBS evaluates the ratio of multipath reflections observed in the out-of-band direction information to prevent erroneous mappings. Further, BBS aggregates the direction estimates retrieved from frames under small scale mobility. As reflections fluctuate under small movements [7], [8] while the direct path does not, this improves the direction mapping precision. BBS further extracts the angular spread of the direct path estimate from the aggregated information. This, gives an indicator for the estimation precision and if multiple mm-Wave antenna sectors come into consideration for mapping. If this is the case, a highly efficient refinement process that uses the already obtained coarse grained direction estimate, determines the optimum sector. The BBS mechanism combines the above techniques to transfer precise and reliable direction estimates to a mm-Wave interface to reduce and in some cases eliminate in-band beam training overhead.

We implement the key components of BBS on a multi-band testbed that combines 2.4 GHz direction inference based on the software defined radio platform WARP [9] and a mechanically steerable 60 GHz RF-frontend used to obtain signal strength measurements. Our evaluation targets three key aspects of BBS. First, combining out-of-band direction inference results and mm-Wave received signal strength analysis, we evaluate mapping precision for multiple locations. We find BBS achieves excellent direction estimation accuracy for short range line-of-sight links that are typical for IEEE 802.11ad. Second, we evaluate the suitability of the multipath reflection ratio to predict a reflected or blocked path and find that it is a very good indicator for both adverse conditions. Lastly, we analyze the trade-off between overhead reduction and achievable maximum receive power for BBS trained directional links. Our analysis shows that while achieving optimum sector selection, BBS eliminates more than 80% of IEEE 802.11ad beamforming overhead.

II. SYSTEM ARCHITECTURE

In this section we describe the BBS architecture and its integration with IEEE 802.11ad. We first explain how directional communication is used in IEEE 802.11ad and then introduce the BBS architecture that combines multiple radio interfaces at different frequencies for out-of-band direction inference.

A. IEEE 802.11ad and mm-Wave WiFi

Directional Communication. To reach the necessary link budget for Gbps throughput mm-Wave communication requires high-gain directional antenna arrays or switched directional antennas. In contrast to 2.4/5 GHz frequencies, a higher amount of mm-Wave antenna elements fits into the same space due to the smaller wave length, which yields higher directional gain [2], [10]. Usage of directional antenna solutions however requires transmitter and receiver to be aligned [1], [2]. In case of mm-Wave communication the propagation behavior is quasi-optical, due to high free space attenuation and limited reflections [11]. Thus, alignment of antenna patterns equals a geometrical matching of transmit and receive focus on the direct geometrical path. IEEE 802.11ad implements code book like predefined beam patterns, with each pattern focusing signal energy into a specific azimuthal direction. With higher number of antenna elements, higher gain and sharper focus of sectors can be achieved. This however leads to a higher number of sectors necessary to cover the full azimuth and the IEEE 802.11ad standard permits up to 128 sectors per device [3]. The number of sender-receiver sector pairs that needs to be matched multiplied by the number directional links to be trained in a network defines the amount of necessary beam training in a directional network.

Establishment of Directional Links. The establishment of directional links or beamforming (BF) training in IEEE 802.11ad is a process required by a pair of nodes to obtain the necessary link budget to perform communication. Through bidirectional BF signaling, each node resolves antenna settings for both transmission and reception [3], [4].

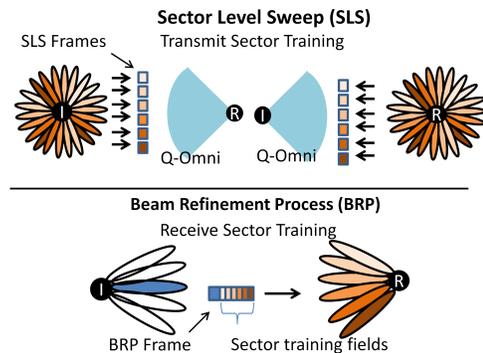


Fig. 1: Beamforming training in 802.11ad.

BF training in the 802.11ad standard consists of two stages as illustrated in Fig. 1: Sector Level Sweep (SLS) and an optional Beam Refinement Process (BRP). The SLS phase enables communication between two nodes at a low PHY rate. In this phase, the best coarse-grain transmit sector is selected for both initiator and responder nodes [3], [4]. To select a transmit sector, an initiator transmits a probe frame on each sector. The responder node receives these frames in a quasi-omni directional antenna configuration. This process is inverted to train the opposite node's transmit sector. The SLS phase is a high overhead procedure, where a complete frame has to be transmitted at every sector at the lowest PHY rate. BBS completely removes the SLS phase from directional link establishment. Instead of performing exhaustive sector search, our system utilizes out-of-band direction information to select appropriate antenna sectors.

In the BRP phase, antenna settings found during SLS using quasi-omni reception are fine-tuned. Also, receive antenna training is added to achieve highest directional gain and multi-Gbps throughput [3], [4]. In contrast to the SLS, the BRP phase can already rely on an established directional link and sweeping of antenna configurations can happen throughout one frame. Thus, a BRP refinement adds significantly less overhead to beam training than a SLS. The BRP phase can be a complementary procedure to BBS when the provided direction information by BBS needs further refinement. However, as shown in our evaluation (Section IV-E) our system reduces the BRP search space and further reduces overhead.

60 GHz for Mobile Devices. Due to sensitivity to blockage and directional antenna misalignment, as well as high attenuation, first generation commercial mm-Wave systems targeted short-range ($< 10m$) quasi-static scenarios, such as HDMI cable replacement. The recently ratified standard IEEE 802.11ad goes the next step and targets the mobile device in a WiFi like application scenario [3], [6]. Also, recent works claim that tenfold the distance of current systems is achievable and even propose mm-Wave cellular systems [5]. However, mobility still poses a problem, as overhead for directional link setup is amplified due to beam readjustment. We found for a beam width of 7 degree that a mismatch of 18 degree reduces signal strength by 17 dB. Comparing to the IEEE 802.11ad modulation sensitivities [3], it is likely that a link breaks if

rate adaptation and beam recovery mechanism do not react quickly. Considering a walking user 5 meters away from an access point, which sways a mobile device with 70 degree per second, such a mismatch occurs roughly 5 times a second. BBS specifically targets applications with mobile devices and proposes a novel device architecture to obtain target device bearing without in-band overhead cost.

B. Node and System Architecture

Node Architecture. The BBS node architecture is based on a multi-band capable device design where mm-Wave and IEEE 802.11ac/n interfaces are combined. We expect IEEE 802.11ad devices to comply with this design as the standard foresees a session transfer feature between mm-Wave and legacy bands [3]. Further, with the recent acquisition of Wilocity by Qualcomm, we see chip vendors trending towards multi-band chipsets (2.4/5/60 GHz). For the BBS system design we assume a mm-Wave interface that comprises an antenna array with predefined highly directional sectors patterns that cover an 360° azimuth. We refer to the mm-Wave system as the *Application Band*, where the antenna pattern is steered according to BBS information to achieve multi-Gbps directional transmissions. We further assume an IEEE 802.11ac/n interface that has a N-antenna omni-directional array usable for inference of bearing to a pairing node. We refer to this frequency band (802.11ac/n platform) as the *Detection Band*.

Devices applying high directional gain require more precise direction estimation for antenna section selection. However, even devices with few detection band antennas can help to guide high gain application bands. BBS applies a final sector refinement when the direction inference is not accurate enough. A low number of detection band antennas increases the search space for this refinement but still prevents the high overhead sector level sweep phase. Even when equipped with a single antenna in the detection band, a device can still be traced by a multi-antenna BBS detection device, increasing the link training efficiency at one side of the link.

System Architecture. Fig. 2 depicts the BBS system architecture. The IEEE 802.11ad compliant *application band* is on the top part of Fig. 2. At the application band, beam-forming overhead is removed by applying direction estimates provided by BBS, shown in the middle block. BBS obtains preamble sample frames from the *detection band* (depicted in the bottom block). It infers the pairing node's bearing from frames received at lower frequencies (2.4/5 GHz), taking into consideration the antenna geometry of the detection band.

BBS runs a background process that infers other devices' direction by passively overhearing detection band frames. This comes without further overhead and does not require any changes to the detection band protocol. As discussed in section IV-F, even idle nodes generate sufficient traffic (e.g. through power management or AP roaming) for at least coarse grained direction inference. While BBS is independent of the technique used to infer a pairing device's bearing, it requires information about the received signal strength in relation to

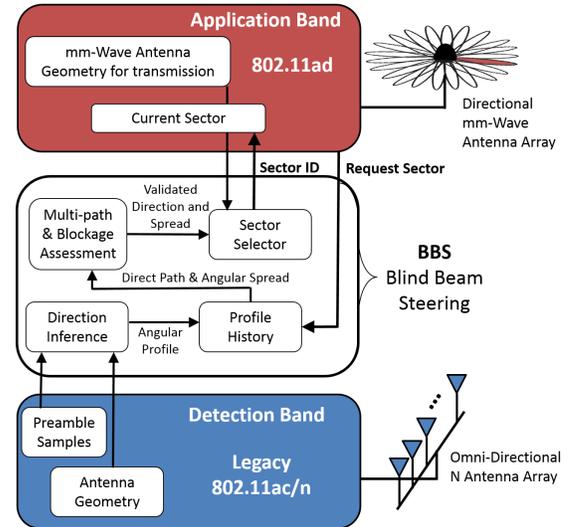


Fig. 2: BBS system architecture.

the azimuthal incidence angle θ . We refer to this information as an angular profile $P(\theta)$. These angular profiles are further analyzed to ensure robustness to mobility, multipath and signal blockage, to then perform reliable mm-Wave sector selection. An example for angular profiles can be found in Fig. 3a. The BBS architecture gathers angular profiles in a history, that is queried whenever a beam training request occurs.

III. MECHANISM

BBS integrates into an IEEE 802.11ad implementation at the point of beam training. Both, link initialization as well as link retraining profit from BBS sector selection. The mechanism in Algorithm 1 ensures that BBS delivers reliable direction information from the detection to the application band. Starting from a beam training request, the angular profile collecting background process (cp. Section III-A) is queried (line 1). If available, a profile history from frames received under differing multipath conditions is returned for the beam training partner. Profiles are combined to increase resilience against multipath (line 3). The aggregated profile is further analyzed for a blocked direct path or remaining multipath effects (line 4). If neither are detected (line 5), the strongest incidence angle is assumed to match the direct path to the training partner and is mapped to application band sectors. As direction estimation uncertainty due to noise is taken into consideration, it is possible that multiple sectors fall into the range of the estimated direct path. To select among these, a low overhead stand-alone BRP is performed (cp. line 9). The selected sector is then returned to the IEEE 802.11ad implementation. In case BBS can not provide a sector estimation, legacy beam training is initiated. The remainder of this section explains the details of Algorithm 1.

A. Out-of-Band Sector Inference and Profile History

BBS performs out-of-band direction inference, using passively overheard detection band frames to calculate angular profiles $P(\theta)$. A profile specifies received signal energy with

respect to the azimuthal incidence angle θ (cp. Section II-B). The retrieved profiles are organized in a history for every overheard device d as defined in Eq. 1. Hereby, $P_t(\theta)$ is an angular profile obtained from a frame overheard at time t and $s(P)$ denotes the node id for the device that transmitted the frame. Every pair of profiles in $H(d)$ is spaced by more than the channel coherence time T_c . This ensures the multipath characteristics to vary between all profiles. Without this additional constraint, the frame aggregation process described in the following subsection, could be biased negatively towards a certain reflected path.

$$H(d) = \{P_t(\theta) | s(P_t(\theta)) = d\} \quad (1)$$

$$\forall P_t, P_v \in H(d) : P_t \neq P_v, |t - v| > T_c$$

Further, the background process ensures that outdated profiles, retrieved from obsolete transmitter positions are removed from the history in a first in first out buffer like process. To this aim, a time weighted average over the profile's main direction component can be used to identify the last accurate frame and crop the history accordingly. Weighting by time, prioritizes newer frames as they are more likely to reflect the current position of the target. We find from our experiments, that a small number of frame preamble samples is sufficient to generate a reliable profile. Thus, even acknowledgment or null data frames can be used to retrieve profiles.

B. Profile History Aggregation

When receiving a beam training request BBS queries the profile history $H(d)$ matching the beam training target d . This history contains profiles obtained under varying multipath conditions and from the latest known position of the pairing node, as described in Section III-A. For the history we found the following two conditions to hold:

- An unblocked LOS path is reflected in every profile by a peak at the same angle.
- Peaks resulting from reflections vary among profiles.

We exploit this by averaging the profiles over the history and define the aggregated angle profile as follows.

$$A_d(\theta) = \frac{\sum_{P \in H(d)} P(\theta)}{|H(d)|} \quad (2)$$

By aggregating over measurements taken under varying multipath conditions, alternating reflection peaks are flattened and the remaining strongest peak likely corresponds to the direct path. A second effect that we observe is that noise and multipath affected frames slightly deviate the direct path angle. Thus, spectra aggregation spreads the direct path peak according to the amount of noise. The spreading gives an estimate of the uncertainty for the direct path estimate.

C. Line-Of Sight Inference and Reflected Path Rejection

When applying direction estimates from the detection band to the application band we have to deal with two major obstacles. First, mm-Wave communication suffers extreme signal attenuation from direct path blockage [11]. This poses a problem when the detection band uses lower frequencies that

Algorithm 1: BBS Sector Selection

Input: beam training request to pairing node d
Result: application band sector at θ^*

- 1 $H(d) = \{P_t(\theta) | s(P_t(\theta)) = d\}$
 $\forall P_t, P_v \in H(d) : P_t \neq P_v, |t - v| > T_c$
- 2 **if** $|H(d)| > 0$ **then**
- 3 $A_d(\theta) = \frac{\sum_{P \in H(d)} P(\theta)}{|H(d)|}$
- 4 $\Psi(A_d) = \frac{\max_{\theta} (A_d(\theta))}{\frac{1}{|A_d|} \sum_{\phi=0}^{2\pi} A_d(\phi)}$
- 5 **if** $\Psi(A_d) > T_{\Psi}$ **then**
- 6 $\theta^* = \operatorname{argmax}_{\theta} (A_d(\theta))$
- 7 $W_{\theta^*} = \operatorname{argmin}_x \left(\frac{A_d(\theta^* \pm x)}{A_d(\theta^*)} \leq T_{\text{peak}} \right)$
- 8 **if** $2 \cdot W_{\theta^*} > W_s$ **then**
- 9 $\text{BRP}(\theta^* \pm W_{\theta^*})$
- 10 **end**
- 11 **else**
- 12 $\theta^* = \text{legacy beam training}$
- 13 **end**
- 14 **else**
- 15 $\theta^* = \text{legacy beam training}$
- 16 **end**

are less prone to blockage. In this case mapping a direction estimate can steer the application band antenna focus into an obstacle, severely impacting link quality and throughput. Second, multipath can induce destructive interference to the direct path signal, in case reflected and direct path signal can not be resolved with sufficient precision. For systems with limited detection band resources this can result in reflected signal strength to be higher than that of the direct path, and the direction estimate to deviate from the direct path. By averaging over multiple profiles as described in Section III-B this effect can be mitigated but not fully prevented.

Reflected mm-Wave paths impose additional attenuation [11] and reflections on the detection band do not necessarily coincide with those on the application band. We therefore do not use reflection based direction estimates for BBS and the system reverts to the high overhead but more resilient legacy beam training method. The same accounts for direct path blockage. To identify these conditions, passively collected detection band information is used and thus no further protocol overhead is imposed.

Peak to Average Ratio. The basic indicator for reflected and blocked path rejection is the ratio of the highest signal strength component to the average received signal energy of an aggregated profile.

$$\Psi(A_d) = \frac{\max_{\theta} (A_d(\theta))}{\frac{1}{|A_d|} \sum_{\phi=0}^{2\pi} A_d(\phi)} \quad (3)$$

Strong multipath spreads the incidence energy over a wide range of the angular receive profile and thus increases the denominator of the fraction. A reflection-less direct path on the other hand results in a sharp peak at its incidence angle and increases the ratio. We experimentally find direction blockage

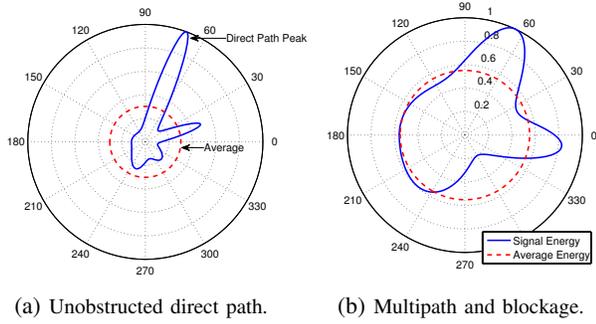


Fig. 3: Angular profile peak to average ratio.

to cause the same effect. As the direct path peak loses magnitude by blockage, the numerator of the ratio shrinks while the average energy received through reflections remains. Fig. 3 shows examples for high and low peak-to-average ratios. We experimentally evaluate a threshold T_Ψ for the peak to average ratio that allows reflected path and blocked line-of-sight rejection at the same time. If the ratio for a direction estimate is below this threshold, BBS proceeds with the legacy IEEE 802.11ad beam training method.

D. Sector Mapping

To relate a direct path estimate from the detection band to the application band, we determine the strongest signal peak identified by its azimuthal angle $\theta^* = \operatorname{argmax}_\theta(A_d(\theta))$ from the aggregated profile for the pairing node d . BBS translates θ^* to the application-band sector that matches the direct path between two nodes and provides the highest signal quality. Due to quasi-optical propagation behavior and strong directional focus of mm-Wave antennas, a mapping is a straight forward geometrical matching. If transmit and receive antenna geometry differ, geometrical matching can be done separately.

E. Optional Sector Refinement

Due to remaining noise and multipath effects, the strongest peak θ^* can slightly deviate from the direct path to the target node. Thus, we extract the peak width around θ^* as it gives insight about the angular range that the direct path can reside in. We define the width of the strongest peak as the minimum angular range around θ^* before it drops below a relative attenuation threshold T_{peak} . The width of the strongest peak W_{θ^*} is defined by Eq. 4.

$$W_{\theta^*} = \operatorname{argmin}_x \left(\frac{A_d(\theta^* \pm x)}{A_d(\theta^*)} \leq T_{\text{peak}} \right) \quad (4)$$

Choosing a low value for T_{peak} extends the angular region around θ^* and increases the chance for the direct path to lie within it. On the other hand, chances for mapping of the estimate to a single application band sector decrease with the size of the considered region. We evaluate the impact of T_{peak} by measurements on our BBS prototype, applying both, a conservative strategy maximizing the likelihood of choosing the optimum sector and a overhead minimizing strategy.

When the angular spread exceeds the width of an application band antenna sector W_s , additional in-band refinement is

triggered with the aim of finding the optimum sector in the indicated set. As a rough estimate for the direction is given by BBS, the IEEE 802.11ad SLS can be skipped and efficient refinement in form of a stand alone BRP phase is carried out. In contrast to the SLS phase that transmits one frame per sector, during BRP a set of sectors is swept with one frame that has a matching number of training fields appended (compare Fig. 1). Thus, the set of sectors reported by BBS can be refined with just one additional in-band frame that sweeps all of them. We experimentally evaluate the overhead reduction achieved by BBS, finding that applying a standalone BRP is a highly efficient method to refine the remaining set of antenna sectors.

IV. IMPLEMENTATION AND EVALUATION

In this section we describe our implementation, the testbed setup, and the performance evaluation of the BBS system.

A. BBS Prototype

We prototype BBS using the FPGA-based software defined radio platform WARP for direction inference in the detection band at 2.4 GHz, and combine this system with a Vubiq 60 GHz waveguide development system [12] for directional communication in the application band.

1) *Implementation of the Detection Band:* For our experiments, we use WARPLab [9], a programming environment capable of prototyping physical layer algorithms in MATLAB that permits over-the-air channel characterization with WARP nodes. For out-of-band sector inference, we use eigenvector analysis based on the MUSIC algorithm [13]. To collect data for the calculation of angular profiles, we use two WARP boards with 8 transceiver chains with synchronized baseband and radio clocks. The setup is shown in Fig. 4. We implement a linear array geometry and refine angle profiles by applying spatial smoothing and antenna sub-grouping mechanisms proposed in [7] with grouping factor $N_G = 2$. Through spatial smoothing the angular profiles are averaged over multiple antenna sub-groups. This further reduces the likelihood of destructive interference on the direct path through adverse multipath conditions in certain locations. As sub-dividing a 4 antenna configuration leaves our profile calculation with 3 sample streams a maximum of two reflections can be resolved. Thus, to evaluate the performance for devices with few detection band antennas we disable the spatial smoothing mechanism. Further, the spectrum history length is fixed to a value of 50 IEEE 802.11g frames. All angular profiles obtained for one location comply with the requirements for the profile history described in III-A. To generate the angular profile, for every frame three short preamble sequences (corresponding to 192 samples) are averaged. Further, the direct path estimate for every receiver location and its mean error are averaged over 500 histories with 50 frames each. The tracked device is a WARP board with one radio interface.

2) *Implementation of the Application Band:* We implement the application band for directional multi-Gbps data transmission using a mm-Wave development platform from Vubiq [12]. This platform consists of transmitter and receiver waveguide

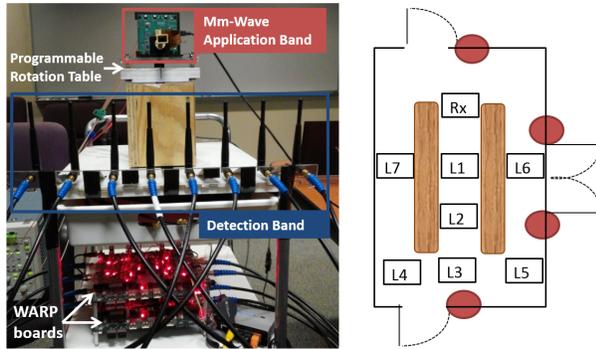


Fig. 4: BBS experimental platform and measurement setup.

modules that provide 1.8 GHz modulation bandwidth in the 58 to 64 GHz unlicensed channels. We feed a 16-QAM modulated multi-tone baseband signal with 15 MHz bandwidth from an Agilent E4432B signal generator to the mm-Wave transmitter. We analyze the received signal strength with a Tektronix TDS7054 oscilloscope measuring the IQ baseband amplitude provided by the mm-Wave receiver. To evaluate different sector patterns we utilize four types of antennas: three horn antennas with beamwidths of 7, 20 and 80 degrees, and an omni-directional antenna. We implement an IEEE 802.11ad like sector sweep through a programmable rotating table that emulates fixed sector positions. The implementation of the mm-Wave application band is shown in Fig. 4.

3) *Experimental Setup*: We conduct our experiments in an indoor setting. To match the targeted mm-Wave WiFi scenario, we perform single-room experiments. Our experiments are conducted under static channel conditions in a room of 15x8 meters, where we perform measurements in 7 locations as shown in Fig. 4. All transceivers are placed at a height of 1-1.5 meters and non-obstructed direct path conditions are ensured, except for the experiments analyzing blockage. We take a series of frame preamble snapshots for the 2.4 GHz tracking target capturing 0.4 ms of the transmitted signal, while the delay between snapshots is approximately 2 seconds. We infer direction on different subsets of this data to evaluate the performance of 4, 5, 6, 7, and 8 detection band antennas. As our static measurement environment experiences long channel coherence times we place the tracking target on a rotation device. This induces small position changes that do not significantly change the angle between nodes but alter multipath reflection characteristics.

B. Direct Path Detection Accuracy

The key requirement to achieve multi-Gbps rates with mm-Wave communication is a non-obstructed direct path between transmitter and receiver. It is vital that the sector inferred from the 2.4 GHz detection band corresponds to this direct path, but effects such as reflections and scattering at the 2.4 GHz band affect the detection accuracy. In this section, we evaluate BBS's accuracy to infer the direction to a pairing node and select the correct sector for directional communication.

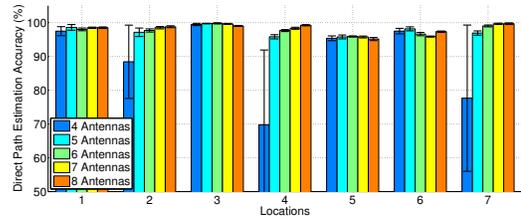


Fig. 5: Detection accuracy of direct path for 8 to 4 detection antennas in 7 evaluated locations.

We determine the detection accuracy for every measured location as the ratio of angular mismatch against the maximum mismatch of 180 degree. The angular mismatch is determined from 60 GHz signal strength measurements and the 2.4 GHz direction estimation results. It is defined as the absolute difference between the direct path's center peak on 2.4 GHz and the center angle of the strongest 7 degree beamwidth antenna sector on 60 GHz. Fig. 5 shows the detection accuracy (Y-axis) for all 7 transmitter locations (X-axis) sub-divided by the number of used detection band antennas.

The detection accuracy results show excellent direct path estimation accuracy with an average precision of 97.8% over all locations when using at least 5 detection band antennas. Accuracy varies slightly over locations indicating varying multipath severeness for different transmitter positions. The number of detection band antennas shows no significant impact on direct path estimation accuracy except for a 4 antenna detection band configuration, for which the direct path prediction accuracy drops to 89%, 67%, and 78% at locations 2, 4 and 7, respectively. The increased sensitivity to multipath for the 4 antenna configuration results from deactivating spatial smoothing as described in section IV-A1. *We find that on average, BBS achieves 97.8% (corresponding to 4 degrees) accuracy in mapping 2.4 GHz direction estimates to mm-Wave antenna sectors with beam widths as low as 7 degrees when using at least 5 detection band antennas.*

C. Robustness to Multipath and Signal Blockage

Blocked direct paths can severely reduce the received signal of mm-Wave links, rendering communication impossible. Furthermore, as discussed in Section IV-B multipath introduces inaccuracies in direct path detection, and might cause the direction estimation to lock onto a reflected path. The following two experiments evaluate BBS's ability to identify these conditions.

Detection of Severe Multipath. In order to analyze the effect of multipath, we evaluate the peak to average ratio as defined in Section III-C. A high ratio indicates little impact by multipath and vice versa. Thus, we compare the accuracy results of IV-B to the ratio to verify that it can be used as an indicator for the accuracy of BBS estimates. Fig. 6 depicts the relation of the direction estimation accuracy achieved by an angular profile to its peak to average ratio. The data points are aggregated over locations and number of detection band antennas. From the graph we observe a clear correlation between high peak to average ratio and accuracy.

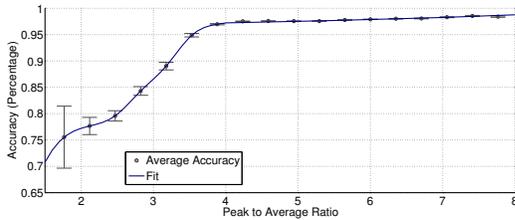


Fig. 6: Peak to average ratio of aggregated profiles in relation to accuracy.

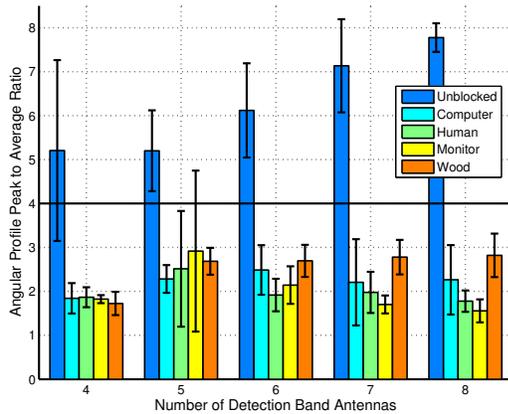


Fig. 7: Effect of blocking obstacles on the mean peak to average ratio of aggregated angular profiles for 7 evaluated locations.

For values below 4, the accuracy decreases significantly with the ratio. Further, the mean error increases with decreasing accuracy which indicates that the error results from locking onto reflection peaks with equal magnitude as the direct path.

We thus define a peak to average threshold T_{Ψ} of 4 and reject aggregated profiles with a ratio below. We evaluate over all locations and antenna configurations the number of false positives and negatives, where a threshold based decision wrongly turns off BBS or uses a highly inaccurate estimate with accuracy below 90%. We observe that a threshold based decision strategy can successfully detect strong multipath and prevent erroneous mapping with a success rate of 94%

Detection of Blocked LOS Path. Our third experiment evaluates the impact of LOS blockage on the spectrum peak to average ratio. This experiment was performed in locations 1 and 3, which have a receiver-transmitter distance of 2 meters and 9 meters respectively. The results are averaged over these locations. The blockage material applied to the direct path is typical to an office environment, i.e., a desktop PC, a monitor, plywood with a width of 1.8 cm, and human blockage.

In Fig. 7 we observe the effect of blocking obstacles in the aggregated angular profile peak to average ratio. The Y-axis depicts the peak to average ratio and the X-axis depicts the number of detection band antennas. The results show that for all types of blocking objects, the peak to average ratios do not exceed a value of 3. When choosing a peak to average ratio threshold of 4 (as proposed from the multipath analysis), we assure with 96.5% certainty that the obtained direction estimate refers to a non-obstructed LOS path. BBS

reliably distinguishes between unblocked and blocked paths by a threshold based evaluation of the peak to average ratio.

D. Training Overhead

BBS maps the estimated direct path to mm-Wave antenna sectors under consideration of the angular spread detected for the direct path peak. When multiple mm-Wave sectors fall into this region, a standalone IEEE 802.11ad BRP selects the optimum among them. The angular spread depends on the relative signal attenuation threshold T_{peak} that sets the boundaries of the peak. Low values for T_{peak} increase the search space and thus the likelihood to identify the best sector, but also increase training overhead.

With higher device complexity, detection and application band antenna orders increase. For example an AP has more 2.4/5 GHz antennas as a smart phone and also implements stronger directional focus on the mm-Wave band. Thus, we assume a correlation between the antenna order on both bands and analyze the behavior of three device classes: 4, 6, and 8 detection band antennas with 80, 20, and 7 degree sector beamwidth in the application band, respectively.

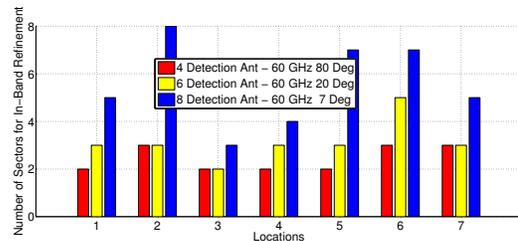


Fig. 8: Overhead: remaining sectors for in-band refinement with optimal sector selection for $T_{\text{peak}} = 0.3$.

Optimum Sector Selection. For our measurement campaign, we find $T_{\text{peak}} = 0.3$ to be the threshold that always includes the optimum sector while having the narrowest possible refinement search space. Thus, BBS always finds the optimum sector and reaches the maximum achievable received power under this parameterization. In this experiment, we evaluate the number of mm-Wave sectors necessary for BRP in-band refinement for this optimum sector selection strategy. We obtain these numbers from the average angular spread W_{θ^*} of each location mapped to the sector width of the three device classes. Fig. 8 depicts on the Y-axis the amount of sectors considered for BRP refinement. The X-axis shows the transmitter locations sub-grouped by BBS antenna configurations. We observe that the BRP search space is increased for complex antenna systems, i.e., 8 detection antennas with 7 degree 60 GHz sectors. A maximum in-band overhead of 7 sectors for such a 52 sector system is required at location 2 which corresponds to 13% of the maximum possible search space. The smallest amount of required in-band refinement can be found in location 3 where only 2 sectors are required for a 52 sector system and no overhead is imposed for the other two antenna configurations. These findings correlate with the high detection accuracy found in IV-B. BBS can achieve

In-Band Sector Refinement	7 Degrees			20 Degrees			80 Degrees		
	IEEE 802.11ad SLS+BRP time = 1.54 ms BBS (ms)	Time Savings (ms)	% Reduction	IEEE 802.11ad SLS+BRP time = 0.88 ms BBS (ms)	Time Savings (ms)	% Reduction	IEEE 802.11ad SLS+BRP time = 0.63 ms BBS (ms)	Time Savings (ms)	% Reduction
0	0	1.54	100	0	0.88	100	0	0.63	100
2	0.28	1.26	81.42	0.28	0.6	67.55	0.28	0.35	54.66
4	0.29	1.25	81.03	0.29	0.59	67.21	0.29	0.34	53.72
10	0.31	1.23	79.87	0.31	0.57	64.85	-	-	-
20	0.34	1.20	77.94	-	-	-	-	-	-

TABLE I: BBS and IEEE 802.11ad time comparison for directional link establishment.

perfect sector selection with only 0 to 13% remaining in-band training.

E. Time of Directional Link Establishment

As shown in Section IV-D, BBS provides perfect sector matching, and thus optimum link budget. In order to achieve the perfect sector matching, there are cases in which BBS requires beam-refinement for a small number of sectors. Nevertheless, BBS removes the IEEE 802.11ad sector level sweep phase and reduces the BRP search space, thus significantly reducing the amount of time required to establish directional links. This directly leads to an increase in spectrum efficiency by allowing more time for data transmissions. In this experiment we evaluate the BBS overhead reduction by comparing the link setup time of IEEE 802.11ad to our proposed method. Specifically, we compare the time BBS takes to perform a standalone BRP on the remaining sector set after out-of-band direction inference compared to the full IEEE 802.11ad legacy beamforming training procedure comprising SLS and BRP.

In case of BBS refinement sector set is determined by the method described in Section III-D. For the legacy method we assume the search space to be fixed to one fourth of the total number of sectors. For the BRP procedure, we assume a three step message exchange for both BBS and IEEE 802.11ad. During the BRP, first initiator and receiver realize a receive training transmitting one frame with the number of requested training fields. Second, two interleaved receive and transmit BRP transactions are performed which include feedback for identified strongest sectors. For all calculations a pairing node with 16 sectors is assumed. Table I lists the BBS and IEEE 802.11ad time comparison for directional link establishment for the evaluated antenna configuration of BBS. We extract the average BRP search space size for the three antenna configurations from Fig.8 (0 sectors, 2 sectors and 4 sectors) and evaluate the corresponding training time reaction. We find that the BBS guided beam forming training achieves an average channel usage reduction of 1.24 ms, 0.59 ms and 0.35 ms per link setup for sector widths of 7, 20 and 80 degrees respectively. The corresponding values are highlighted in Table I. This performance increase scales with the number of nodes in a network and link retraining caused by device mobility. Further, for highly directional systems a notable setup delay reduction is achieved by BBS. For example, to setup a link that needs in-band refinement of 4 sectors, only 0.29 ms are required compared to 1.54 ms for a 7 degree beam width system. *Summing up, BBS achieves perfect sector selection, while reducing beamforming training overhead for every link setup by up to 1.24 ms (81%) for highly directional systems with sector widths as narrow as 7 degrees.*

F. Direct Path Detection under Mobility

Due to the lack of commercially available mm-Wave antenna arrays, that support controllable real-time beam steering, it is difficult to directly evaluate BBS under mobility. Nonetheless, by reducing the frame history length in our static setup we can deduce the behavior for mobile nodes, where direction information become obsolete fast. We observe that the high average accuracy above 95% can be maintained for frame histories as small as 8 frames, when at least 5 detection band antennas are available. Further, we measured the number of frames transmitted by a range of different idle IEEE 802.11n/ac devices (notebooks and smartphones). We found a substantial amount of null data frames to be transmitted and their number to increase with the expected device mobility. This renders coarse-grained direction estimation under mobility possible, given BBS' small frame history length requirements. Null data frames serve for several implementation dependent features e.g. power management, channel scanning, keep alive mechanisms and AP roaming.

Additionally, IEEE 802.11ad defines a seamless session transfer from mm-Wave to legacy frequencies in case of link breakage at the mm-Wave band. This feature creates lower frequency traffic, which allows BBS to derive a high accuracy direction estimate and reestablish the link.

V. RELATED WORK

Prior work on 60 GHz communication mainly focused on channel measurements and modeling, as well as hardware design. To the best of our knowledge, BBS is the first work to introduce multi-band coupling to remove overhead for beam-training in directional mm-Wave communication. However, there exists complementary work on the key components of our system: multi-band coupling and direction inference.

Multi-band Systems. Various previous systems and hardware architectures concurrently use multiple bands for purposes other than directional link establishment. For example, Wake on Wireless [14] uses multiple interfaces on battery operated devices to save energy. It puts devices on standby mode and utilizes a second very low-power radio at a different frequency band to wake up the primary device when required. Other works combine 3G, WiMAX and WiFi radios to increase capacity or offload traffic [15], [16]. Protocols for opportunistic usage of multi-band spectrum were studied in [17]. The 802.11ad standard [3] itself supports seamless integration with 802.11a/b/g/n/ac, through a mechanism called transparent Fast Session Transfer (FST). This mechanism allows real-time transition of communication from any band/channel to any other band/channel. As a consequence, we expect a large

fraction of IEEE 802.11ad transceivers to operate on multiple bands and thus directly support our BBS mechanism.

Direction Inference. The techniques utilized for out-of-band detection inference in the present paper are based on [7], [8], [18]. These mechanisms have primarily evolved from MUSIC [13], which was the first work to analyze incident waveforms: quantity, direction of arrival, strength and cross correlation among them. Direction inference in combination with various multi-antenna or MIMO techniques and channel impulse response measurements is often used for indoor localization [7], [8], [19], [18]. Incidence angle based localization techniques have further been used in ad-hoc mesh networks [20], CDMA mobile cellular systems [21], [22] and in combination with ToA using antenna arrays and interference cancellation for NLOS environments and wide-band CDMA [22]. In contrast, BBS utilizes bearing inference in a multi-band system with the purpose of guiding 60 GHz directional communication.

60 GHz Beamforming Overhead. In [23] beam training is considered as optimization of a 2-D signal strength function defined over finite codebooks of transmit and receive antenna sectors. The proposed algorithm, based on a numerical Rosenbrock search, probes a reduced number of sector combinations compared to the exhaustive searches that are currently state of the art. However, probing specific sector pairs requires coordination between the involved nodes and reliable communication on untrained links is not possible in IEEE 802.11ad networks. Thus, the method proposed in [23] is not applicable to this type of network. In contrast, the IEEE 802.11ad standard [3] uses a sector level sweep to find antenna sectors for initial communication and BBS efficiently solves the problem by out-of-band direction inference.

VI. CONCLUSION

In this paper, we design and experimentally evaluate a novel architecture and framework for mm-Wave directional communication, named *Blind Beam Steering (BBS)*. BBS addresses the overhead problem for directional mm-Wave link establishment that poses a limit in mm-Wave's very high throughput capabilities, beam-width scalability and mobility support. Our system architecture couples 60 GHz mm-Wave with legacy 2.4/5 GHz bands to exploit propagation properties of each. We infer the line-of-sight direction between the communicating devices from omni-directional transmissions at low frequencies in the *Detection Band*. The information is further processed to ensure robustness to multipath and signal blockage, to then perform mm-Wave sector selection. The sector selected by the BBS framework is passed to the *Application Band* which implements IEEE 802.11ad compliant mm-Wave directional communication. Our experimental evaluation demonstrates that BBS achieves on average 97.8% accuracy for direction estimation when 5 or more antennas are used at the detection band. This accuracy in direction estimation results in IEEE 802.11ad beam forming training overhead reduced by 81% for a highly directional transceiver with 7 degree beam width. Further, in our typical office-style testbed

environment BBS successfully detects unblocked direct path conditions for highest throughput mm-Wave transmissions with a rate of 96.5%.

VII. ACKNOWLEDGMENTS

This research was supported by the EU's Seventh Framework Programme ERC Grant Agreement no. 617721, the Keck Foundation, a grant from Intel Corp., and NSF grants CNS-1444056, CNS-1126478 and CNS-1012831.

REFERENCES

- [1] R. C. Daniels and R. W. Heath, "60 GHz Wireless Communications: Emerging Requirements and Design Recommendations," *IEEE Vehicular Technology Magazine*, 2007.
- [2] J. Ning, T.-S. Kim, S. V. Krishnamurthy, and C. Cordeiro, "Directional Neighbor Discovery in 60 GHz Indoor Wireless Networks," in *Proc. MSWiM*, 2009.
- [3] IEEE 802.11 Working Group, "IEEE 802.11ad, Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band," 2012.
- [4] T. Nitsche, C. Cordeiro, A. B. Flores, E. W. Knightly *et al.*, "IEEE 802.11ad: Directional 60 GHz Communication for Multi-Gigabit-per-Second Wi-Fi," *IEEE Communications Magazine*, 2014.
- [5] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson *et al.*, "Demystifying 60GHz Outdoor Picocells," in *Proc. ACM MobiCom*, 2014.
- [6] P. Moorhead, "Qualcomm Gets The Jump On WiGig 60 GHz Wireless With Wilocity Acquisition," July 2014. [Online]. Available: <http://www.forbes.com/sites/patrickmoorhead/2014/07/02/qualcomm-gets-the-jump-on-60-ghz-wireless-with-wilocity-acquisition/>
- [7] J. Xiong and K. Jamieson, "ArrayTrack: A Fine-grained Indoor Location System," in *Proc. NSDI*, 2013.
- [8] —, "Towards Fine-grained Radio-based Indoor Location," in *Proc. ACM HotMobile*, 2012.
- [9] "Rice University WARP project," available at: <http://warp.rice.edu>.
- [10] M. Park, C. Cordeiro, E. Perahia, and L. L. Yang, "Millimeter-Wave Multi-Gigabit WLAN: Challenges and Feasibility," in *Proc. IEEE PIMRC*, 2008.
- [11] H. Xu, V. Kukshya, and T. S. Rappaport, "Spatial and Temporal Characteristics of 60-GHz Indoor Channels," *IEEE Journal on Selected Areas in Communications*, 2002.
- [12] Vubiq, V60WGD03 60 GHz Waveguide Development System. [Online]. Available: <http://vubiq.com/v60wgd03/>
- [13] R. O. Schmidt, "Multiple Emitter Location and Signal Parameter Estimation," *IEEE Transactions on Antennas and Propagation*, 1986.
- [14] E. Shih, P. Bahl, and M. J. Sinclair, "Wake on Wireless: An Event Driven Energy Saving Strategy for Battery Operated Devices," in *Proc. ACM MobiCom*, 2002.
- [15] S. Jindal, A. Jindal, and N. Gupta, "Grouping WI-MAX, 3G and WI-FI for wireless broadband," in *IEEE and IFIP International Conference in Central Asia on Internet*, 2005.
- [16] X. Hou, P. Deshpande, and S. Das, "Moving bits from 3G to metro-scale WiFi for vehicular network access: An integrated transport layer solution," in *Proc. IEEE ICNP*, 2011.
- [17] A. Sabharwal, A. Khoshnevis, and E. W. Knightly, "Opportunistic spectral usage: Bounds and a multi-band CSMA/CA protocol," *IEEE/ACM Transactions on Networking*, 2007.
- [18] J. Gjengset, J. Xiong, G. McPhillips, and K. Jamieson, "Phaser: Enabling Phased Array Signal Processing on Commodity WiFi Access Points," in *Proc. ACM Mobicom*, 2014.
- [19] C. Wong, R. Klukas, and G. Messier, "Using WLAN Infrastructure for Angle-of-Arrival Indoor User Location," in *Proc. IEEE VTC*, 2008.
- [20] D. Niculescu and B. Nath, "Ad hoc Positioning System (APS) Using AOA," in *Proc. IEEE INFOCOM*, 2003.
- [21] L. Xiong, "A Selective Model to Suppress NLOS Signals in Angle-of-Arrival (AOA) Location Estimation," in *Proc. IEEE PIMRC*, 1998.
- [22] L. Cong and W. Zhuang, "Hybrid TDOA/AOA Mobile User Location for Wideband CDMA Cellular Systems," *IEEE Transactions on Wireless Communications*, 2002.
- [23] B. Li, Z. Zhou, H. Zhang, and A. Nallanathan, "Efficient Beamforming Training for 60-GHz Millimeter-Wave Communications: A Novel Numerical Optimization Framework," *IEEE Transactions on Vehicular Technology*, 2014.