

X-Array: Approximating Omnidirectional Millimeter-Wave Coverage Using an Array of Phased Arrays

Song Wang*
University of California, San Diego
sowang@ucsd.edu

Jingqi Huang*
University of California, San Diego
jih032@ucsd.edu

Xinyu Zhang
University of California, San Diego
xyzhang@ucsd.edu

Hyoil Kim
Ulsan National Institute of Science
and Technology
hkim@unist.ac.kr

Sujit Dey
University of California, San Diego
sdey@ucsd.edu

ABSTRACT

Millimeter-wave (mmWave) networks are conventionally considered to bear a fundamental coverage limitation, due to the directional beams and limited field-of-view (FoV) of the phased array antennas. In this paper, we explore an array of phased arrays (APA) architecture, which aggregates co-located phased arrays with complementary FoVs to approximate WiFi-like omni-directional coverage. We found that straightforwardly activating all the arrays may even hamper network performance. To fully exploit the APA's potential, we propose X-Array, which jointly selects the arrays and beams, and applies a dynamic co-phasing mechanism to ensure different arrays' signals enhance each other. X-Array also incorporates a link recovery mechanism to identify alternative arrays/beams that can efficiently recover the link from outage. We have implemented X-Array on a commodity 802.11ad APA radio. Our experiments demonstrate that X-Array can approach omni-directional coverage and maintain high performance in spite of link dynamics.

CCS CONCEPTS

• **Hardware** → **Wireless devices; Analysis and design of emerging devices and systems**; • **Networks** → **Network experimentation; Network mobility**;

KEYWORDS

millimeter-wave networks, mmWave, array of phased arrays, multi-panel phased-array, beam management, 5G

ACM Reference format:

Song Wang, Jingqi Huang, Xinyu Zhang, Hyoil Kim, and Sujit Dey. 2020. X-Array: Approximating Omnidirectional Millimeter-Wave Coverage Using an Array of Phased Arrays. In *Proceedings of The 26th Annual International Conference on Mobile Computing and Networking, London, United Kingdom, September 21–25, 2020 (MobiCom '20)*, 14 pages. <https://doi.org/10.1145/3372224.3380882>

1 INTRODUCTION

The emerging wireless infrastructure is facing a massive mobile

*Both authors contributed equally to this work.

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiCom '20, September 21–25, 2020, London, United Kingdom

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7085-1/20/09...\$15.00

<https://doi.org/10.1145/3372224.3380882>

traffic demand [1–4], driven by billions of upcoming Internet of Things and immersive multimedia applications [5–9]. Due to the spectrum crunch in legacy low-frequency bands, both the wireless local area and cellular network standards have been incorporating millimeter-wave (mmWave) technologies (e.g., 802.11ad and 5G NR) to meet the looming challenge of mobile traffic overload. Ideally, one would expect the mmWave technologies to provide WiFi or LTE-like seamless coverage. But mmWave signals have orders of magnitude higher attenuation loss, which has to be compensated through high-gain phased array antennas. The directional gain of a phased array is, empirically, proportional to the number of antenna elements [10–13]. Hence, to provide sufficient mmWave coverage, an intuitive way is to simply increase the phased array size.

However, the high directionality brings two new challenges to mmWave networks: (i) *Beam management overhead*. A phased array may generate hundreds of beam patterns with main lobes pointing to different directions. Ideally, by rapidly scanning through the beams, it can approximate the behavior of an omni-directional antenna. Yet, when the receiver is mobile or when the line-of-sight (LoS) is blocked by obstacles, such a trial-and-error scanning may incur huge overhead in finding an alternative beam [14–16]. Recent years witnessed substantial research in designing efficient algorithms to identify the optimal beam directions under link dynamics [17–20]. (ii) *Limited field-of-view (FoV) coverage*. A phased array has limited FoV (typically narrower than 120° [21–23]), due to the intrinsic properties of its patch antenna elements. This problem remains largely underexplored. Very recently, a Pia system [24] was proposed to combine multiple access points (APs) to expand the FoV [24], but this requires dense deployment, which is often not economically feasible, and involves tight coordination among distributed APs.

A more viable approach to overcoming the FoV limitation is to aggregate standard phased arrays to form an “array of phased arrays” (APA). In an APA mmWave radio, multiple phased arrays are co-located, sharing the same RF chain but facing different angles to jointly cover 360° in azimuth or elevation. The combined coverage may provide more *multipath diversity*, i.e., signals may traverse different paths through reflections, making it easier to save a link under blockage. The APA architecture has been adopted recently by both 802.11ad [25, 26] and 5G NR devices [27–32]. Apart from coverage, APA also offers advantages in cost and efficiency. This is because beyond certain physical dimensions of the radio package, the feed network losses (between the radio RF front-end and a single giant phased array) would negate the benefits of having higher array gains [21]. In addition, all the phased arrays can share the same codebook, thus reducing the on-chip register/memory requirements, which accounts for non-trivial cost on a radio device [33].

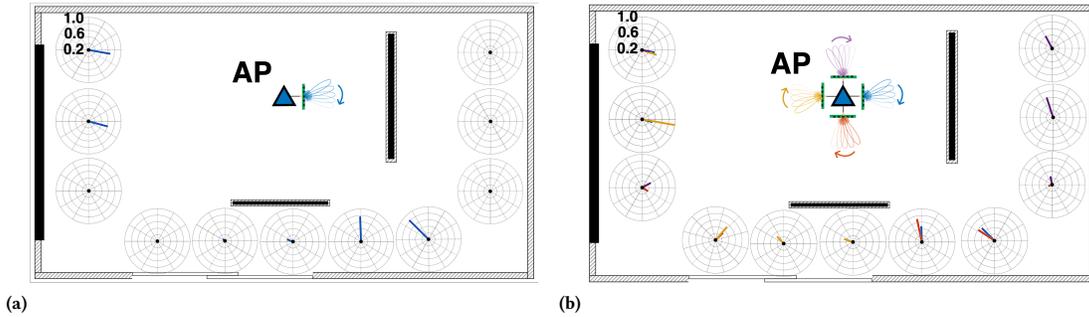


Figure 1: Coverage and multipath diversity under the same power constraint: (a) Single array. (b) 4-array APA.

Overall, APA makes it practical to scale to large phased arrays [34].

Unfortunately, APA also incurs new challenges to mmWave network design. *First*, due to the regulation constraint on emission power, not all phased-arrays can be turned on simultaneously. So a node has to decide on not only which beam to use for each array, but also how many and which arrays to activate. The decision space easily escalates to an intractable scale. To our knowledge, no existing work attacked such a problem of array/beam management. *Second*, the APA should leverage its advantages in multipath diversity, to efficiently recover from link outage caused by mobility, blockage, or a mix of both. This would require it adapt the beam/array selection in real-time, with minimum protocol level overhead. *Third*, not all phased arrays' signals combine in a coherent way, so a straightforward way of turning on multiple arrays may even lead to lower link quality than a single array.

In this paper, we propose a novel system called *X-Array*, which explores the challenges and opportunities from APA, through three major design components. (i) We propose an optimization-driven array/beam selection algorithm to maximize the link quality under the power budget constraints. The solution is formulated as a look-up table, which is generated in an offline one-time manner and maps the dominant signal paths' angle-of-departure (AoD) to the optimal array/beam combination. At run-time, the APA node only needs to run a simple AoD estimation algorithm, leveraging the periodic beacon scanning defined in mmWave standards such as 802.11ad. (ii) We apply a low-overhead *dynamic co-phasing* algorithm to the different transmit arrays, so that their signals can coherently combine at the receiver, with very infrequent feedback. This method can maximize the combined directional gain (under regulation constraint), while maintaining as wide beamwidth as possible, to make the link more resilient under mobility. (iii) We design a link recovery mechanism that leverages the multi-array architecture to efficiently and accurately find alternative arrays/beams when the strong path disappears (e.g., due to blockage) or reappears, with minimal overhead.

We have implemented X-Array on a commodity single RF chain 802.11ad AP supporting up to 8 arrays. X-Array runs at the device's user space, so it does not require any hardware modification. We have conducted experiments to verify X-Array in diverse radio environments, including indoor and outdoor, with different levels of mobility, multipath conditions, and blockage dynamics. The results demonstrate that: (i) X-Array can approach the best array/beam combinations for all the settings, and the corresponding overhead is comparable with single-array solutions. (ii) X-Array can efficiently update the array/beam selection under link dynamics caused by mobility and blockage. (iii) X-Array can correctly apply the co-

phasing factors to the multiple phased arrays to maximize the advantages of APA, while respecting the transmit power constraint. (iv) The joint coverage achieved through X-Array is comparable to an oracle solution that exhaustively searches across the decision space. By dynamically selecting and switching among 8 arrays or even 4 arrays, X-Array can approximately achieve omni-directional coverage in a sophisticated environment with random blockage. In contrast, a naive solution with a single array or two fixed arrays leave many blind spots with extremely low bit-rate.

Although commercial APA 802.11ad radios already exist [25, 26], they typically turn on all arrays which is far from optimal and may even be worse than a single-array (Sec. 5). X-Array represents the first system to fully exploit the advantages of APA. Our main contributions can be summarized as follows. (i) An efficient way to jointly manage multiple phased arrays and their beams to maximize link quality; (ii) A multi-array joint beamforming mechanism to ensure coherent combination of the multiple array's signals; (iii) A link recovery mechanism to ensure the robustness of APA under blockage; (iv) Implementation and validation of X-Array on a commodity 802.11ad APA radio.

Our implementation of X-Array essentially converts a commodity multi-array 802.11ad radio [26] into a partially programmable experimental platform. Unlike recently developed mmWave software radios [35, 36], this platform can only run the 802.11ad MAC/PHY and does not provide channel state information (CSI). But it is less costly, and it allows for selecting beam, codebook, and arrays. Latest development of the platform will be documented in [37], and instructions for using the platform will be provided upon reasonable requests.

2 MOTIVATION AND CHALLENGES

In this section, we conduct preliminary experiments to demonstrate the potential benefits of the APA mmWave radio, and the practical challenges in harvesting the benefits.

2.1 Potential Advantages of APA

We first investigate the unique channel characteristics of APA in comparison with single-array. Our experiments run on an off-the-shelf 802.11ad APA radio. The radio supports multiple arrays, each being a 6×6 uniform planar array with around 120° FoV. More detailed hardware specifications will be introduced in Sec. 4. As a benchmark experiment, we place a single-array AP and a 4-array AP in the middle of a $11\text{m} \times 6\text{m}$ indoor meeting space. For a fair comparison, we *configure the two APs to use the same set of beams on each array, and the same total power constraint*. Two poster boards (with metal plate on the back) stand nearby, representing reflec-

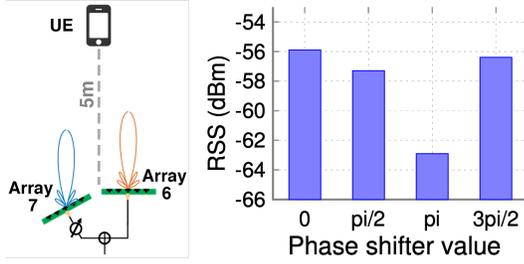


Figure 2: The impact of co-phasing.

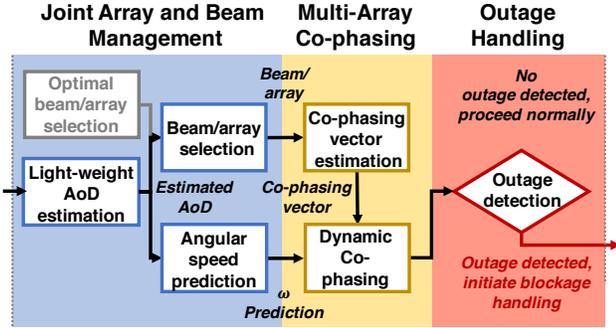


Figure 3: X-Array workflow.

tors/obstacles just as in a typical environment.

During the experiment, a user carries a client device, walking while keeping natural body orientation, so occasionally her own body blocks the LoS of the AP-client link. We measure the *AoA profile* at the client side, at 11 random locations. The *AoA profile* depicts the received signal strength (RSS) along each angular direction. We omit weak AoAs that are 5 dB lower than the strongest one, because of their negligible contributions to the total RSS. From a high level, the AoA profile shows the multi-path diversity and each path’s quality at each specific location. For simplicity, we only use APA at the AP side, whereas the client is tuned to a quasi-omni beam, but the effects can be reciprocal. Fig. 1 (a) and (b) show the resulting AoA profiles. Each line segment in the polar plot represents the AoA of a signal path and the segment length denotes the corresponding RSS. We have two major observations from the results.

Limitation of single-array: A single-array AP has very limited FoV and creates very limited multipath. As shown in Fig. 1 (a), only a few locations within the single array’s FoV have reasonable RSS. Beyond those are the AP’s blind spots and the client can only rely on the NLoS reflected signals, which tend to be weaker and come up in a sporadic way. Overall, the single-array AP can only provide very few dominant paths for only locations within its FoV. We note that recent work [24] characterized the FoV constraint of 802.11ad radios, but the focus was on the antenna gain pattern, so reflection/blockage effects are not analyzed.

The benefit and potential of APA: An APA is able to expand the coverage dramatically and provide more multipath diversity, potentially leading to more robust links under mobility and blockage. With a 4-array AP, more locations have strong LoS paths owing to the complementary FoVs of multiple arrays, even for those originally in the blind spots of the single-array case. In addition, almost all client locations receive signals from multiple dominant AoAs, *i.e.*, the APA can make better use of surrounding reflectors, since the enlarged FoV contains signals with more diverse AoAs. This benefit exists

even for locations whose LOS are blocked. The additional multipath will be valuable for maintaining a robust connection, because the signals can be resteeered along a new path’s direction even if one is blocked. However, note that for certain client locations (*e.g.*, top left), the APA leads to weaker RSS than the single-AP. This implies that *it may not be optimal to activate all arrays simultaneously*, since the transmit power is spread out, and certain arrays’ signals may cancel each other.

2.2 Challenges

The foregoing measurement reveals the potential of APA, assuming an oracle system that can orchestrate the arrays and beams with no overhead. In practice, approaching this ideal entails non-trivial design challenges.

Joint array and beam management for mobile users: An 802.11ad AP executes a sector level sweeping (SLS) periodically at the beginning of each beacon interval (BI), where it broadcasts signals sequentially through each of its N beam patterns. The periodic beam scanning ensures the AP is discoverable by unassociated clients, and the best beam can be identified for each associated client. The SLS beam scanning involves N 52-byte beamforming (BF) frames and $N - 1$ 1- μ s Short Beamforming Inter Frame Spacing (SBIFS), two 208-bit SSW-feedback frame and one Long Beamforming Inter Frame Spacing (LBIFS) [38]. The total scanning overhead is relatively small, *e.g.*, ~ 1.1 ms for a 64-beam array, in contrast to the typical BI of 100 ms. However, for mobile clients, the *beam coherence time*, *i.e.*, mean period within which the best beam index remains unchanged, becomes much shorter. So the beam scanning has to be executed more frequently. Specific to the APA, the scanning overhead is further multiplied by the number of arrays, rising to ~ 8.8 ms even for an 8-array AP. Besides, even if we ignore the protocol overhead, the AP has to jointly decide beam selection and array selection (*e.g.*, billions possibilities, Sec. 3.3), a combinatorial problem that can easily exhaust its computational power.

At first blush, one can simply turn on all phased arrays to circumvent the array selection problem. Unfortunately, the FCC regulation [39] imposes constraints on both the total radiation power (TRP) and effective isotropic radiation power (EIRP). The EIRP constraint limits the phased array gain along the peak direction to 43 dBm and the average of all directions to 40 dBm, for safety and interference concerns. The TRP constraint further limits the total emission power of all directions to 500 mW [39]. When all the arrays on an APA are active, the TRP needs to be split among all of them, whereas an optimal solution should concentrate all the power towards the strongest eigen mode of the channel, *i.e.*, beaming the signals towards the strongest AoA.

Co-phasing between phased arrays: The phased arrays on the same APA node share the same RF chain and transmit the same digital baseband signals. However, due to their relative location/orientation differences, and hardware-induced initial phase offset, the emitted signals do not necessarily combine coherently at the receiver. To showcase this phenomenon, we run a controlled experiment with two arrays (index 6 and 7) on the AP facing the client direction. Fig. 2 shows up to 7 dB of variation in RSS, as the relative phase between the two arrays varies between 0 , $\frac{\pi}{2}$, π , and $\frac{3\pi}{2}$. This implies a strong need for phase compensation, or *co-phasing*, to ensure coherent signal combination between concurrent arrays. The problem becomes more pronounced in mobile scenarios, as slight location variation causes significant phase change (due to the short wavelength). The need to choose the optimal co-phasing factor essentially adds one more dimension in the APA’s decision space,

making it intractable. Note that our 802.11ad radio only allows configuring the co-phasing factor with a 2-bit resolution (4 relative phase values). Finer phase resolution will further compound the decision complexity.

Recovering from link outage, especially under blockage: MmWave link outage may occur in an unpredictable manner, due to other objects moving across the LoS or the device user's own body blockage. For single-array radios, existing work has explored algorithms to realign the transmitter and receiver's beams, taking advantage of the correlation between beam patterns on the same phased array [17, 19]. Yet for an APA, a new mechanism is needed that can reselect the array as well as its optimal beam to leverage the multipath diversity. And again, a brute-force way of rescanning all arrays may incur non-trivial overhead, especially when the user is moving and body blockage occurs frequently.

3 DESIGN

3.1 Design Overview

We now briefly introduce X-Array's design components and workflow. We build X-Array on top of the 802.11ad MAC/PHY stack. Without loss of generality, we assume a single client served by the X-Array access point. Extension to multiple users can be straightforwardly realized using the built-in MAC protocol in 802.11ad, *i.e.*, transmitting to each client sequentially with CSMA/TDMA based scheduling. We assume the AP uses an APA whereas the mobile client has a single phased array due to form-factor constraint. To ensure it is discoverable by clients facing arbitrary directions, the X-Array AP has to follow the 802.11ad SLS (Sec. 2) to periodically broadcast a beacon frame through each beam and repeat it for each array. Considering the overhead of such *full scanning* (Sec. 2), it has to be activated infrequently (default to every 8 BIs in X-Array).

X-Array's main design components and decision logic run on the multi-array AP, shown in Fig. 3. Whenever a client is associated, the AP runs a one-time full-scanning. The client calls a lightweight *AoD estimation algorithm* (Sec.3.3) and feeds back its estimation to the AP. Given the current AoD, the AP uses a lookup table to select the optimal array(s) and beams to activate. The table only needs to be generated once in an offline manner, using a *joint array/beam selection algorithm* (Sec.3.3) that optimizes the overall APA beam output pattern with respect to each AoD, under the TRP and EIRP constraints. Afterward, the client and AP proceed to their runtime routine, and periodically update the array/beam selection, based solely on the per-BI SLS beacon broadcast from one of the activated arrays.

Meanwhile, whenever two or more arrays are activated, X-Array applies a *multi-array co-phasing algorithm* (Sec.3.4) to ensure that the arrays' signals are coherently combined. To align the signal phases, X-Array AP estimates the inter-array phase difference, and then compensates the difference by applying an initial phase offset (*i.e.*, co-phasing factor), which is allowable in commercial 802.11ad APA hardware (Sec. 4). As will be verified in Sec. 3.4, keeping track of the inter-array phase offset directly incurs huge measurement overhead. Therefore, the X-Array AP predicts the phase offset changes within one BI instead, based on the changing rate of the estimated AoD. It continuously applies the co-phasing factor based on the predicted phase offset, until the beginning of the next BI when the AoD is refreshed.

Occasionally, the mmWave link may experience an outage, *i.e.*, low or null RSS, likely due to AP losing track of client under blockage coupled with abrupt motion. Then the AP executes a novel *multi-array concurrent beam scanning* scheme to rediscover strong

signal paths and reidentify the best array/beam. This scheme reduces the outage recovery overhead from 8.8 ms to around 1.2 ms on an 8-array node (64 beams per array), compared with a full scan.

3.2 Preliminaries: Modeling APA Multi-Array Beamforming

We first introduce a model of APA, which is a basis for the exposition of the X-Array design. For simplicity, we assume a Uniform Linear Array (ULA) with antenna elements arranged along the azimuth plane with half-wavelength displacement. Note that the design can be easily extended to Uniform Planer Array (UPA), and our implementation uses 6×6 UPA.

Modeling single-array. For a single ULA with N antenna elements, assuming omni-directional Rx, the received signal can be formulated as:

$$\mathbf{y} = \mathbf{H}^T \mathbf{w} \mathbf{x} + \mathbf{n} \quad (1)$$

where \mathbf{H} and \mathbf{w} both are 1-by- N vectors, representing the channel gain from the N transmit antennas, and the beamforming weights, respectively. \mathbf{x} represents the transmitted symbol and \mathbf{n} represents the noise.

The separation between antenna elements, d , is usually half wavelength, much shorter than link distance, so it suffices to model the far field. The channel can be decomposed as gain component \mathbf{A}_G , which can be approximated to be consistent across antenna elements, and phase component $[e^{j\frac{2\pi nd}{\lambda} \sin(\phi)}]^N$, $n = 0, 1, \dots, N-1$, where ϕ is Angle-of-Departure(AoD), the angle between the normal line of the phased array panel and receiver wave-front. Plugging this decomposition in Eq. (1), we have:

$$\mathbf{y} = \mathbf{w}^T \mathbf{A}_G [e^{j\frac{2\pi nd}{\lambda} \sin(\phi)}]^N \mathbf{x} + \mathbf{n} \quad (2)$$

where $[\cdot]^N$ represents a vector of size N , and $n = 1, 2, \dots, N$. The phase component $[e^{j\frac{2\pi nd}{\lambda} \sin(\phi)}]^N$ is usually called *steering vector* for phased array beamforming. To maximize the directionality gain towards ϕ , the complex conjugation of the steering vector at ϕ should be used as the *codebook entry* to form a *beam*, *i.e.*,

$$\mathbf{w}_\phi = [e^{-j\frac{2\pi nd}{\lambda} \sin(\phi)}]^N \quad (3)$$

Each codebook entry constitutes one row in the *codebook* (a matrix). The codebook entries are often designed to steer to R angles that equally partition the FoV. Note that X-Array is applicable to other codebook design objectives as well (*e.g.*, minimizing sidelobes [40]).

Modeling APA. Now consider an APA transmitter with P phased arrays. Similar to Eq. (1), we have:

$$\mathbf{y} = \mathcal{H}^T \mathbf{w} \mathbf{x} + \mathbf{n} \quad (4)$$

where \mathbf{w} and \mathcal{H} are a 1-by- NP beamforming weight vector and a 1-by- NP channel gain vector, for all NP antenna elements on the APA, which can be seen as a new, giant phased array. *Since the phased arrays may be placed in different positions/orientations, not all NP antenna elements follow a ULA layout. So the beamforming equation Eq. (3) does not necessarily hold for APA.* Alternatively, we can regard the Rx signal as the coherent combination of signals from the Tx arrays. Then Eq. (4) can be rewritten as:

$$\mathbf{y} = \sum_{p=0}^{P-1} \mathbf{w}_p^T H_p \mathbf{x} + \mathbf{n} \quad (5)$$

where H_p is 1-by- N channel gain vector of the phased array with index p . Since the two phased arrays are co-located on the same

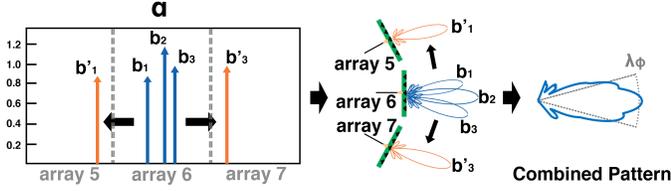


Figure 4: X-Array optimization relaxation: replace and redistribute beams to match λ_ϕ .

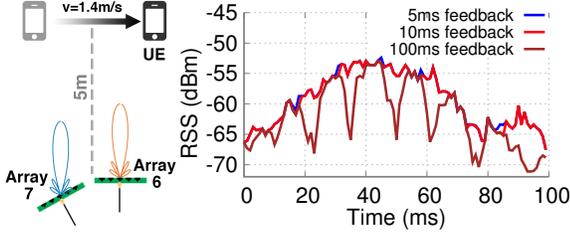


Figure 5: Mobility causes two arrays to lose co-phasing periodically, unless with frequent feedback.

device, and typically only separated by several centimeters, shorter than the Tx-Rx distance. Thus, we can again make the far-field assumption, *i.e.*,

$$\mathbf{y} = \sum_{p=0}^{P-1} \mathbf{w}_p^T \mathbf{H}_p e^{j \frac{2\pi s_p}{\lambda} \cos(\phi + \frac{\delta_p}{2})} \mathbf{x} + \mathbf{n} \quad (6)$$

where s_p is the displacement between the centroids of array p and a reference array, and δ_p denotes the angle between their normal directions. \mathbf{H}_p is the channel gain vector for array p , similar to the term \mathbf{H} in Eq. (1). We define the inter-array phase difference as *co-phasing vector* $\mathbf{E}_\phi = [e^{j \frac{2\pi s_p}{\lambda} \cos(\phi + \frac{\delta_p}{2})}]^P, p = 1, 2, \dots, P$.

3.3 Joint Array and Beam Management

3.3.1 Optimal Array/Beam Selection for a Given AoD. We formulate the joint array/beam selection as an offline optimization problem. For a given AoD, the objective is to appropriate the optimal beam pattern to maximize the power towards the AoD direction subject to TRP and EIRP constraints, while creating maximum multipath diversity. Note that the one-time offline optimization solely depends on the line-of-sight AoD. Hence it does not need to rerun when environment changes.

Without loss of generality, we assume the APA jointly covers an FoV of 360° on the azimuth plane where the clients are located. We equally partition the 2π FoV into R directions and denote each partition as $\varphi_r = \frac{2\pi(r-1)}{R}, r = 1, 2, \dots, R$. The vector of RSS values of all R directions, or normally called “beam pattern”, of the beam indexed b on array p , can be represented as:

$$\mathbf{r}_{b,p} = |\mathbf{w}_{b,p}^T [e^{j \frac{2\pi n d}{\lambda} \sin(\varphi_r)}]^{N \times R}| \quad (7)$$

where $[\cdot]^{N \times R}$ represents a N -by- R matrix; $n = 1, 2, \dots, N$ and $r = 1, 2, \dots, R$. Here we omit the channel gain factor A_G in Eq. 2 as it contributes equally for all beams. Correspondingly we express the collection of beam patterns on a phased array p as $\mathbf{R}_p = \{\mathbf{r}_{b,p}, b \in B\}$. As we mentioned earlier, the phased arrays in X-Array share the same codebook. Hence *the beams of the same index on different arrays share the same beam patterns, although they might point at different directions due to the arrays’ orientation differences.*

The optimization problem can be formulated as:

$$\max_{\alpha} \lambda_\phi \sum_{b,p} \alpha_{b,p} \mathbf{r}_{b,p}^T \quad (8)$$

$$\text{s.t.} \quad \sum_{b,p \in \mathcal{E}} \alpha_{b,p} \mathbf{r}_{b,p} \leq I_{\text{EIRP}} \text{ (elementwise)} \quad (9)$$

$$\| \sum_{b,p} \alpha_{b,p} \mathbf{r}_{b,p} \|_1 \leq I_{\text{TRP}} \quad (10)$$

$$\alpha_{b,p} \in \{0, 1\}, \forall b, p \quad (11)$$

$$\sum_b \alpha_{b,p} \leq 1 \quad (12)$$

where λ_ϕ is a 1-by- R vector associated with a given AoD ϕ , and α is a B -by- R binary decision matrix. The λ_ϕ is the ideal APA beam pattern for the given AoD ϕ . For simplicity, we define λ_ϕ as a simple binary vector which has unit gain within $\pm 20^\circ$ of its intended AoD and 0 gain elsewhere. Eq. (10) and (9) are the TRP and EIRP constraints, respectively. The constraints (11) and (12) represent the fact that we can choose at most one beam on one array. λ_ϕ represents a customized ideal beam pattern. The goal of the maximization function is to find a selection of beams whose combined pattern best matches λ_ϕ . We define λ_ϕ as a simple binary vector where the direction $\pm 15^\circ$ around AoD has 1 and others 0. We show in Sec. 5 that even with such simple definition, X-Array can substantially improve the beam coherent time (Sec. 2.2).

The maximization objective (8) essentially combines the RSS values of individual beam patterns as the resulting joint beam pattern of multiple beams on multiple arrays. In other words, we model the average power combination of the beams. Later we will introduce the dynamic co-phasing design (Sec. 3.4) which ensures the different beams’ signals are coherently combined at run time to further enhance SNR.

This optimization framework is non-linear, and the search space is determined by the number of beams B and arrays P . A brute-force way of solving the problem requires searching across B^P beam patterns, *i.e.*, 8 arrays with 64 beams on each requires $64^8 \approx 2.8 \times 10^{14}$ times computation which is intractable. Thus, we relax the constraints (11) and (12). That is, we allow choosing an arbitrary number of beams on one array. Through this relaxation, we transform the previous combinatorial optimization into a standard linear optimization problem:

$$\max_{\alpha} \lambda_\phi \sum_{b,p} \alpha_{b,p} \mathbf{r}_{b,p}^T + \beta \sum_{b,p} \alpha_{b,p} \quad (13)$$

$$\text{s.t.} \quad \sum_{b,p \in \mathcal{E}} \alpha_{b,p} \mathbf{r}_{b,p} \leq I_{\text{EIRP}} \text{ (elementwise)} \quad (14)$$

$$\| \sum_{b,p} \alpha_{b,p} \mathbf{r}_{b,p} \|_1 \leq I_{\text{TRP}} \quad (15)$$

$$\alpha_{b,p} \geq 0, \forall b, p \quad (16)$$

Here we also add the sum of α to the maximization goal in order to encourage the α to have fewer terms. This relaxed version of array/beam optimization can be solved efficiently by standard linear programming toolboxes.

By relaxation, we allow having multiple beams in one array in α which is infeasible in practice. To fix this problem, we leverage a key observation: The FoVs of multiple arrays are usually partially overlapped, so certain beams on different arrays share similar directions. This indicates we can replace the multiple beams on one array with beams of similar directions on adjacent arrays. As a result, to make

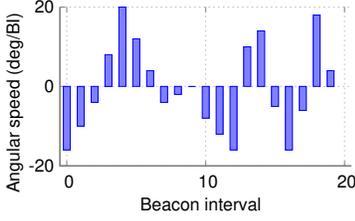


Figure 6: Client angular speed.

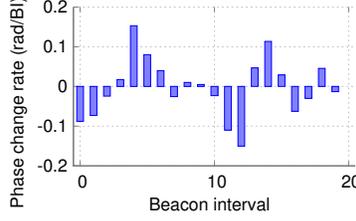


Figure 7: Phase changing rate.

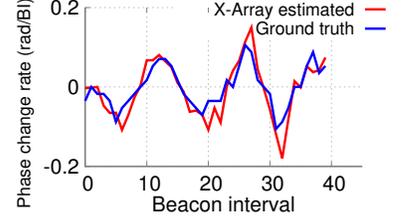


Figure 8: X-Array phase prediction matches the ground truth well.

α feasible, we can simply identify the multiple beams on one array in α , and replace the extra beams with beams of similar directions on neighboring arrays. Fig. 4 illustrates this process. The similarity of directions between arrays is subject to vendor implementation, *i.e.* the size and geometrical layout of phased arrays. Without loss of generality, we use beam AoD to measure the similarity of beams instead of arrays.

In case when there are not enough similar beams to replace, we discard the beams whose directions deviate the most from the AoD, to guarantee the gain along the AoD direction. Note that there may be two reasons for the lack of similar beams: (i) The optimization includes too many non-zeros terms in α . (ii) The arrays do not have much FoV overlap, so the number of similar beams is small. We prevent (i) by adding the regulation term $\beta \sum_{b,p} \alpha_{b,p}$ in Eq. (13). On the other hand, (ii) can be avoided with more arrays (*e.g.*, when per-array FoV is 120°).

3.3.2 AoD Estimation and Tracking. AoD is assumed as a given input in the aforementioned array/beam selection algorithm. To estimate the AoD, X-Array runs the 802.11ad SLS which scans the beams on one of its currently active arrays. Intuitively, we choose a *primary array* whose orientation is closest to previous AoD, and let the client measure the sequence of *per beam* RSS, denoted as $[r_b]^B$ where b is beam index. According to Eq. (2):

$$r_b = |\mathbf{w}_b^T [e^{j\frac{2\pi nd}{\lambda} \sin(\phi)}]^N + \mathbf{n}| \quad (17)$$

Since the SLS beam scanning is very short (Sec. 2), the AoD can be assumed relatively stable during SLS. *Our objective here is to estimate the AoD ϕ based on the $[r_b]^B$ measurement.* Since the codebook weights \mathbf{w}_b and element-spacing d are known, we can compute the $[r_b]^B$ for any given ϕ when omitting noise. The $[r_b]^B$ that best correlates with the measurement should correspond to the most likely ϕ . We formalize this intuition through a matched filter design.

The matched filter is a B -by- R matrix derived from the product of two parts: (i) $[\mathbf{w}_b]^{B \times N}$ is the beam weights of B beams arranged in B -by- N matrix; (ii) $[e^{j\frac{2\pi nd}{\lambda} \sin(\phi_r)}]^{N \times R}$ is normalized steering vector of R directions arranged in N -by- R . As an extension to Eq. (17), the dot product of (i) and (ii) produces the RSS sequences for all R possible AoDs. Then the matched filter simply correlates itself with the measured RSS sequence $[r_b]^B$ as follows:

$$\nabla = [r_b]^B [\mathbf{w}_b]^{B \times N} ([e^{j\frac{2\pi nd}{\lambda} \sin(\phi_r)}]^{N \times R})^T \quad (18)$$

Then we can estimate AoD by identifying the direction r that leads to the maximum similarity between measurement and model:

$$\phi = \arg \max_{r \in R} \nabla \quad (19)$$

X-Array's AoD estimation mechanism leverages the legacy SLS beam scanning in 802.11ad (Sec. 2), so it requires no hardware modification and shares the same overhead (*e.g.*, ~ 1.1 ms out of each BI of around 100 ms). Many existing systems [18, 19, 41] also need AoD

as an input, but they either lack support on commodity hardware (due to needs for CSI), or they require non-trivial computational time.

It is tempting to think that one can detour the array selection problem, by treating the APA as a *single giant phased array* and apply a single codebook to it. However, this *single giant array approach* lacks scalability for two reasons: (i) *Hardware constraints.* The single codebook needs to specify all possible array/beam combinations, which easily reaches billion scale as mentioned above, way beyond the storage capability of on-board memory (only several hundred KB on a typical 802.11 device[33]). (ii) *Protocol overhead.* Scanning through all the entries on the giant codebook takes 4.8×10^9 seconds for the typical 8-array APA, and will obviously hinder normal data transmissions.

3.4 Multi-Array Co-Phasing

X-Array's co-phasing design aims to maximize the beamforming gain when multiple arrays are activated. To overcome the challenge of phase sensitivity (Sec. 2.2), we propose a novel dynamic co-phasing scheme that approximates the optimal coherent combination of multiple arrays, without the overhead of constantly probing their phase offsets.

3.4.1 Decomposition. Recall the multi-array channel can be modeled as a composition of signals from individual arrays through a shared channel with phase offsets (Eq. (6)). To ensure coherent signal combining, the beamforming weights \mathbf{w}^T must be designed to compensate for the different arrays' phase offsets. Based on Eq. (6), we thus have:

$$\mathbf{w}^T = \sum_{p=0}^{P-1} \mathbf{w}_p^T e^{-j\frac{2\pi sp}{\lambda} \cos(\phi + \frac{\delta_p}{2})} \quad (20)$$

Recall $\mathbf{E}_\phi = [e^{j\frac{2\pi sp}{\lambda} \cos(\phi + \frac{\delta_p}{2})}]^P$, $p = 1, 2, \dots, P$, represents the inter-array phase differences, or *co-phasing vector*. We can then rewrite the APA beam weights equation Eq. (20) as:

$$\mathbf{w}^T = \mathbf{w}_p^T \mathbf{E}_\phi^* \quad (21)$$

where $(\cdot)^*$ is the complex conjugate operator. This implies that *the multi-array co-phasing problem can be decomposed as beamforming on individual arrays, but with proper inter-array phase alignment.* Therefore, to realize co-phasing, we do not need to modify the existing codebook. Instead, we can simply multiply each individual array's codebook entry with an offset $[\mathbf{E}_\phi^*]_p$, which is allowable on commodity hardware (Sec. 4). We now describe how to estimate the \mathbf{E}_ϕ .

3.4.2 Estimating Co-Phasing Vector. We adopt a measurement driven method to estimate the \mathbf{E}_ϕ at the beginning of a BI. Specifi-

cally, after each individual array's beam is determined (represented by \mathbf{w}_p^T) on each array p , we regard one of the arrays as *reference array* with phase 0. To measure the relative phase of other currently active arrays (denoted as *side arrays*), the AP transmits 4 BF frames (the reference signal used in 802.11ad) using the reference array and one side array *simultaneously*. We apply a phase shift $e^{j\varphi}$ to each of the BF frames. Specific to our 802.11ad radio, $\varphi \in \{1, j, -1, -j\}$ (Sec. 4). The RSS value for these BF frames can be formulated as:

$$y_\varphi = |\mathbf{w}_p^T \mathbf{H}_p e^{j\frac{2\pi s_p}{\lambda} \cos(\phi + \frac{\delta_p}{2})} e^{j\varphi} + \mathbf{w}_0^T \mathbf{H}_p|, \varphi \in \{1, j, -1, -j\}$$

By applying the phase shift value, we essentially build a discrete Fourier series with $e^{j\varphi}$ as "frequency basis" and the co-phasing vector as "coefficients". Hence, to extract the co-phasing vector, we only need to take a Fourier transform on this series:

$$\mathbf{d}^{|\Phi|} = \text{FFT}(\{y_1, y_j, y_{-1}, y_{-j}\}) \quad (22)$$

Then we find the second term $d_{p,2} = e^{j\frac{2\pi s_p}{\lambda} \cos(\phi + \frac{\delta_p}{2})}$. Recall that the co-phasing vector (Sec. 3.2) is the inter-array phase difference between reference array and another side array. For array p , this phase offset is exactly $d_{p,2}$. Hence we repeat this process for all currently active arrays and we have co-phasing vector:

$$\mathbf{E}_\phi = [d_{p,2}]^P \quad (23)$$

We need 4 BF frames (each lasting 0.015 ms) for every active array except the reference array. Hence, the co-phasing vector measurement for one array takes 0.06ms. Even with 8 active arrays, the total overhead is negligible (< 0.5 ms).

3.4.3 Dynamic Co-Phasing. The foregoing co-phasing algorithm assumes that the AoD information is always available as input. However, the inter-array phase offset varies drastically over a few ms under node mobility, whereas the AoD can only be updated per BI in order to tame the estimation overhead (Sec. 3.3).

More specifically, recall that the steering vector in Eq. (2) has a changing rate of $\frac{2\pi nd}{\lambda} \sin(\phi)$ and the co-phasing vector, as shown in Eq. (6), has a changing rate of $\frac{2\pi s_p}{\lambda} \cos(\phi + \frac{\delta_p}{2})$. The array displacement s_p is significantly larger than d , the antenna element spacing. This implies that, with the same angular movement ϕ of Rx, the APA Tx (affected by both fast-changing co-phasing vector and steering vector) suffers more from link degradation than the single array Tx (affected only by mild steering vector).

Intuitively, one can keep track of the array steering term by probing the AoD and update the co-phasing factor more frequently. However, the small coherence time of the array steering term requires an impractically high feedback frequency to prevent link degradation. As a showcase, we activate two beams on two arrays on an AP, and move the client at walking speed. We repeat the experiment with different phase feedback intervals (100 ms, 20 ms, 5 ms). Fig. 5 shows that the throughput converges only when feedback interval is shorter than 10 ms. Larger intervals cause sub-optimal throughput "valleys" due to laggy feedback. At higher moving speed, even more frequent feedback is needed.

We address this challenge using an *angular speed based dynamic co-phasing* scheme. *Dynamic co-phasing* obtains the fine-grained phase estimates within the scope of one BI, *i.e.* between two consecutive phase feedbacks. At the beginning of a BI, the AP measures the co-phasing vectors for the currently active arrays as the initial co-phasing vector \mathbf{E}_{ϕ, t_0} . Within each BI duration, to avoid the feedback, we predict the optimal instantaneous co-phasing vector by modeling the relationship between the array steering vector and angular speed. Denote $\Delta\tau$ as AoD estimation interval, at a given

time t_j when AoD is updated, we can estimate the average angular speed of the client as:

$$\bar{\omega}_j = \frac{\phi(t_j) - \phi(t_j - \Delta t_\tau)}{\Delta\tau} \quad (24)$$

We assume the angular speed is stable within $\Delta\tau$ since a typical BI is very short. Thus, we can predict the phase change within one $\Delta\tau$ interval as a function of time:

$$\frac{2\pi s_p}{\lambda} \cos(\bar{\omega}_j \tau + \phi(t_j) + \frac{\delta_p}{2}) \approx \text{ang}(\mathbf{E}_{\phi, t_0}) + \bar{\omega}_j \tau \quad (25)$$

To demonstrate the effectiveness of dynamic co-phasing, we leverage the same experimental setup as in Fig. 5, and keep measuring the co-phasing vector every 5ms to get fine-grained ground truth. Meanwhile, we estimate AoD every 100 ms (one BI) and calculate client angular speed by Eq. (24). Fig. 6 and Fig. 7 show the client angular speed and the phase change rate of the co-phasing vector \mathbf{E}_ϕ . The strong resemblance of these two figures further corroborates Eq. (25). We then input the angular speed to our dynamic co-phasing model, and predict the co-phasing vector change over every 5 ms within the next BI. Fig. 8 shows that the predicted phase changing rate matches the ground truth near perfectly. With instantaneous co-phasing vector known, the AP now can align the phase of arrays within a BI without any explicit client feedback.

3.5 Recovering from Link Outage

When blockage occurs, if the LoS still delivers strong RSS or there exists any strong reflection path within the FoV of active array(s), the array/beam management and co-phasing solutions are still applicable (based on the periodic AoD estimation on the primary array). But if a link outage occurs, *i.e.*, weak or null RSS on the current link, then X-Array invokes its outage handling mechanism. A straightforward way is to repeat beam-sweeping on all arrays and select the strongest beam, but apparently this will incur huge overhead.

To tame the rescanning overhead, we propose a simple *concurrent beam sweeping* scheme. Immediately upon outage, X-Array concurrently beacons a reference frame through the same beam index on all arrays, and repeats this for each beam index within the codebook (shared by all arrays). If any strong LoS/NLoS path exists, then at least one beam will lead to a strong AoD peak. If the strong path falls in the FoVs of multiple arrays, then each such array will have one beam with similar RSS as the peak. But the beam indices tend to differ due to the arrays' different orientations (*i.e.*, the same beam index on different arrays points to different directions).

To showcase this phenomenon, we fix an 8-array AP and put a client 3m away, and then snapshot the per-beam RSS of concurrent beam sweeping when a human blocks the LoS path, and after the blockage exits. Fig. 9 (a) and (b) plot the results. It can be seen that, upon blockage, two weak peaks exist on the per-beam RSS sequence, likely because a certain beam can establish a NLoS path. Whereas after blockage disappears, three strong peaks reappear.

If the peak RSS after concurrent scanning still falls below the threshold for the minimum bit-rate, then no array/beam can sustain the link, and extraneous connectivity solution may be needed (*e.g.*, [42]). Otherwise, upon confirming the existence of a usable beam, X-Array needs to further discriminate which array(s) cover the LoS within its FoV. The concurrent scanning result already indicates the best beam indices that lead to strong AoD. So X-Array simply sends a reference frame through the corresponding beam index

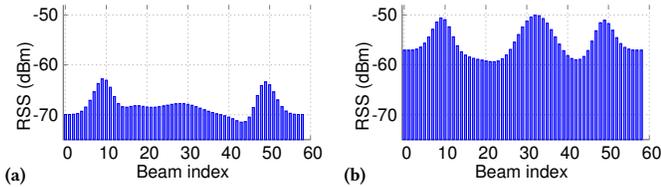


Figure 9: Concurrent beam sweeping when: (a) blockage occurs; (b) blockage disappears.

on each array. Those arrays and beams that lead to AoD peaks will be selected as active arrays. Afterwards, X-Array moves out of the outage mode and transitions into its normal mode of operation. Overall, the concurrent scanning mechanism can be called on to reidentify a strong beam when the current link’s RSS drops significantly. When a link is under blockage, it can also be called periodically to check whether the blockage disappears and a new strong path reappears.

Two additional issues are remarkable here: (i) A straightforward full-scan needs to probe NB beams in total, vs. $(N + B)$ with concurrent beam sweeping. On an 8-array APA with 64 beams per array, this means the latter reduces the overhead of rediscovering strong paths by $\sim 8\times$ (8.8 ms vs. 1.2 ms). (ii) Although co-phasing may sometimes weaken the strongest peak on the RSS sequence, it rarely removes the peak, because the strengthening effects may show up on other beams pointing close to the AoD. Plus, we only need to know *whether* a strong peak exists, so a coarse grained per-beam RSS sequence suffices.

4 IMPLEMENTATION AND EXPERIMENT SETUP

4.1 Implementation

We implement X-Array based on an off-the-shelf 802.11ad AP from Airside Inc. [26]. The original Airside AP puts 8 phased arrays on a plane. We reorganize its physical layout, and 3D print an antenna stand (Fig. 10) so that the 8 arrays face different azimuth directions with 45° separation. In this way, the FoVs of different arrays partially overlap, and together they cover 360° azimuth and 120° elevation. The mobile client has a single phased array. Both devices comprise an 802.11ad NIC (with Qualcomm QCA9500 FullMAC WiGig chip and QCA6335 baseband) with 2-bit phase shifters, plugged in an embedded Linux host (running Wil6210 firmware and driver). Below we describe the notable technical thrusts in implementing X-Array.

(1) *Fine-grained per beam RSS extraction*: The per beam RSS at client side is crucial for the AoD estimation and co-phasing mechanism, but is concealed to upper layers on commodity 802.11ad devices. To expose the RSS, we first disassemble the firmware file to ARC assembly code. Then we blanket search the assembly code and pinpoint where the per beam RSS calculation takes place, and the associated memory address in NIC. Then we leverage Talon-tools [43], a C-based firmware patching framework adapted for the 802.11 radios, and write the firmware patch to copy the RSS value, immediately after a RSS value is calculated, from its original address to a designated memory address at the very back of NIC memory, which can be safely accessed by the host driver. With the patched firmware loaded to the NIC, we then write a Python program to call the `mem_dump` command on the host, and dump the RSS value from NIC to user space.

(2) *Enabling short BI*: Recall that 802.11ad performs SLS beam

sweeping per BI, so a shorter BI may make the beam selection more responsive but at larger overhead. X-Array does not require short BI thanks to the dynamic co-phasing design (Sec. 3.4). However, to obtain the fine-grained ground truth phase measurement (Sec. 5), we need BI as short as possible. The standard 802.11ad radio limits the smallest BI to be 20 ms, which is not enough for this purpose (c.f. Sec.3.4). To overcome this barrier, we follow similar steps as above, to disassemble the firmware and pin-point the BI value memory address. Then we hard code the BI value to the firmware patch. Extremely small BI will lead to inaccurate RSS measurement or even firmware crash. We empirically found 5 ms to be the smallest safe value.

(3) *Real-time codebook loading and inter-array co-phasing*: We implement X-Array’s dynamic co-phasing by loading selected beams with different phase shifts after AoD estimation, and selecting the optimal phase shift based on the dynamic co-phasing design. The 802.11ad standard limits the codebook size (maximum number of beams) to 128. Yet this X-Array implementation requires more than 128 beams (with different phase shift combinations) due to the APA setup. Hence it is necessary to load the codebook at run-time. Normally the codebook file is only loaded from user space to NIC when the interface boots. We use the `HWD_RFC_WRITE_SECTOR` command (0x900, ut_subtype_id: 0x514) in Qualcomm wil6210 driver to write the new codebook to NIC. Then we call `WMI_SET_RF_SECTOR_PARAMS_CMDID` command (0x9A1), which is originally designed to change one entry in codebook, but it can also trigger codebook reloading onto the phased arrays. This way we can change the codebook without rebooting the NIC.

(4) *Enforcing TRP/EIRP constraints*. We enforce the TRP/EIRP constraint by regulating the beams, which is a common practice by COTS phased array devices. Specifically, we enforce the constraints when optimizing the overall APA beam pattern in respect of AoD, as shown in Eq.9 and Eq10. We then solve this optimization problem for each AoD, so that the optimization output beam pattern complies with the EIRP and TRP constraints. In our real-time implementation, we use the pre-calculated beam and array combination, which automatically enforces the constraints.

(5) *Implementation of other components*. We implement the X-Array AoD estimation algorithm on the client side, where the per-beam RSS measurements are performed. The client then feeds back the estimated AoD along with estimated angular speed to the AP in a single 802.11ad packet. The AP acts accordingly, by calling other design components (implemented as python modules on the user space). Thanks to the lightweight X-Array design, the algorithms can run in *real-time* on the embedded PCs of both the AP and client.

4.2 Experimental setup

For comparison, we have also implemented the following APA solutions [41, 44, 45] as baselines.

(i) ACO [41]: We implemented the phase estimation and AoD estimation algorithms proposed in ACO [41]. Just like X-Array, ACO takes per-beam RSS as input, but it can estimate the phase difference between one reference antenna element and all other antenna elements, from which it obtains the CSI. The corresponding AoD is obtained by running the MUSIC [46] over the CSI. To implement ACO, we generate a custom codebook file with two antennas activated for each beam, corresponding to the reference antenna element and the to-be-measured antenna element. We choose the antenna element with index 0 in the codebook file as reference. We

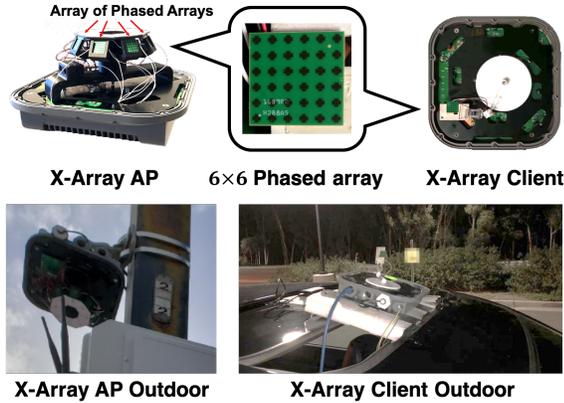


Figure 10: X-Array hardware prototype is built on a commercial multi-array 802.11ad AP, with customized array layout.

generate 4 beams for each to-be-measured antenna with its phase index as 0, 1, 2 and 3 (mapped to $0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}$). We then load the codebook and feed the per beam RSS to the ACO model [41] which is implemented in python.

(ii) Periodic probing: We also implemented a periodic probing mechanism to realize co-phasing across arrays. A probing frequency of $\frac{1}{2}$ BI means that co-phasing is done twice per BI. Here the inter-array co-phasing is then estimated based on the ACO’s CSI output. Higher feedback frequency may make the co-phasing more accurate under mobility/blockage, at the cost of higher overhead.

(iii) Neighbor scan: Since there exists no other work in outage recovery with APA, we implement a *neighbor scan* (NS) baseline. Whenever an outage occurs, NS first scans adjacent arrays with the smallest angle displacement to the previous AoD. If a strong beam exists, it settles on these arrays; otherwise it keeps trying others.

Note that the above *periodic probing*, and *neighbour scan* essentially represent the default behaviors of the existing 802.11 ad protocol when running on an APA radio.

5 EVALUATION

We conduct extensive experiments in three types of environment settings: indoor open area ($16.6m \times 6.4m$ yoga room) with no furniture blockage, complex indoor (typical office environment with workbenches and partitions around the route) and outdoor (parking lot). We will verify that X-Array achieves WiFi-like coverage and maintains high performance under link dynamics. Our results can be summarized as follows:

- X-Array maintains high accuracy in AoD estimation across different environments, enabling a negligible throughput gap (around 5%) in comparison with an ideal solution that knows the ground-truth AoD. Also, X-Array’s joint beam-array optimization can form wider beams without sacrificing the gain, so it can provide stable and high throughput under high mobility scenarios.
- When taking into account co-phasing, X-Array can more accurately predict the phase change of the client within next beacon interval. Consequently, X-Array leads to a much lower throughput gap (<12%) in comparison to periodic probing (18% to 42%).
- The blockage recovering mechanism of X-Array saves most of the link outages and experiences no throughput gap in 96% and 93% of the blockage cases, in complex indoor and indoor open space respectively.
- We put all the components together and run X-Array, and we

find that X-Array does not suffer any throughput gap most of the time under different speeds, while periodic feedback mechanism has a significant throughput loss even with a high frequency of every $\frac{1}{3}$ beacon interval.

- By maximizing the benefits from APA, X-Array can achieve WiFi-like omni-directional coverage. In comparison, a straightforward way of turning on 4 phased arrays may lead to even worse coverage compared with 2 arrays.

5.1 Micro-benchmarks

5.1.1 Joint Array and Beam Management.

Accuracy of AoD estimation. We validate the accuracy and effectiveness of X-Array’s AoD estimation in three different environments. In each environment, we fix the AP and randomly place the client to 80 locations. We compare three schemes: ACO [41], X-Array with and without match filter (the latter referred to as “beam scan only”). The box plot in Fig. 11 shows that X-Array has the lowest average estimation error and lowest std., *i.e.*, it achieves the most accurate and stable AoD estimation. Although ACO performs slightly better in indoor open space, its estimation error increases dramatically in complex indoor and outdoor scenarios, since its CSI estimation only works reliably under high SNR and degrades a lot in relatively long-range and multipath-rich conditions. Also note that the beam scan only approach performs worse than X-Array in all cases due to the imperfect beam patterns, which implies the effectiveness and necessity of the matched filter design (Sec. 3.3).

Impact of AoD estimation error. To understand the end-effect on throughput, we fix the AoD estimation error to a specific value. To control the error, we first measure the groundtruth AoD using a laser range finder. Then we intentionally use an AoD value that deviates from the ground-truth by 2° to 45° , as input to X-Array’s array/beam selection. Note that the commodity 802.11ad device does not allow data transmission under RSS monitoring mode. Moreover, it does not implement high bit-rate 802.11ad modulation and coding, so the benefit of higher channel/link quality cannot be reflected in measured throughput. We thus follow the same approach as in [17, 41] to map the RSS to achievable throughput.

Fig. 12 plots the percentage of throughput loss compared to the ground-truth, denoted as *Normalized throughput gap*. An AoD error of below 5° causes minor throughput loss (*e.g.*, median 10% and 75-percentile at 15% in 4° case), yet the median throughput loss escalates to about 30% for AoD error above 30° . Considering the AoD estimation accuracy (Fig. 11), the corresponding average throughput loss of X-Array falls below 5%, in comparison to the 20%-40% loss of ACO, in complex indoor and outdoor environment.

Effectiveness of X-Array’s joint array/beam selection. We compare X-Array with two baseline approaches, one is to simply lets two arrays beamform to the same direction to emulate the effect of treating the APA as a single giant phased array mentioned in Sec. 3.3, denoted as *w/o optimization*, and the other one is to place multiple arrays with non-overlapping FoV, denoted as *non-overlap*. To verify X-Array’s resilience under mobility, we use *beam coherence time* (Sec. 2) as the performance metric. To create different client moving speeds, we use a time-lapse approach as in [17]: We move the client 2.2cm each time along a 10m trajectory, and collect per-beam RSS traces at each point. Then, different speeds correspond to different time-lapse values between the measurement points. The result in Fig. 13 (a) shows that, even at 75 mph moving speed, the beam coherence time of X-Array remains within one BI (100 ms) and non-overlap remains 1.3 BI (130 ms), whereas the

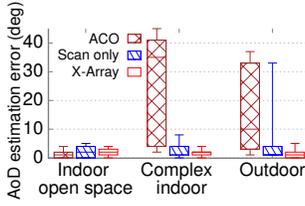


Figure 11: AoD estimation error.

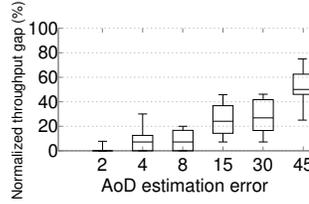


Figure 12: Impact of AoD estimation error.

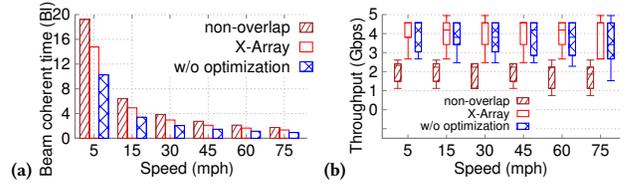


Figure 13: Impact of joint array/beam selection on: (a) link stability; (b) link quality.

giant array can only support up to 45 mph. The reason is that X-Array creates wider beams without sacrificing directionality. Thus, it needs to update the array/beam selection much less frequently, leading to even lower overhead than the giant array approach.

To ensure this benefit does not come at the cost of lower beam quality, we place the client to random positions along trajectory 10 m away, and compare the *average throughput* of the two beam-forming mechanisms. The result is shown in Fig. 13 (b). Fig. 13 (b) shows that, compared with w/o optimization, beams formed by X-Array can achieve comparable or higher average throughput, and lower variance in general. Interestingly, the throughput of both X-Array and w/o optimization is $\sim 2\times$ higher than non-overlap, which indicates that high link quality can be achieved by applying the co-phasing algorithm.

5.1.2 Multi-Array Co-Phasing.

Accuracy of phase prediction. We conduct the co-phasing experiments by arbitrary walking along 10 routes inside an office building, and repeat the aforementioned pointwise measurements to emulate different driving speed in outdoor open space. For each beam combination (with two different beam indices on two arrays) at each measurement point, we exhaustively vary the 4 possible co-phasing values between two arrays, and measure the per-beam RSS when both arrays are activated. As we previously observed in Fig. 8, the phase change rate estimated by X-Array shares a highly similar pattern with the groundtruth. According to the predicted phase change rate, we can predict the two arrays' phase offset across the duration of one BI. The box plot in Fig. 14 further shows the percentile errors when we run phase prediction in indoor (3 mph walking speed) and outdoor (varying speed from 15 mph to 75 mph) settings. We see that X-Array has an average phase error of 0.57 and 0.44 in outdoor and indoor, which are around $2.5\times$ and $5\times$ lower than periodic probing. Moreover, periodic probing incurs around $2\times$ larger phase error in outdoor scenarios than indoor, as it is unable to track the phase change under high speed. In contrast, X-Array becomes even more accurate when it comes to outdoor. The reason is that although the phase change is fast, it is also stable and thus easier to predict.

Impact of phase prediction error. We now evaluate the impact of phase error in terms of the throughput loss compared with an oracle solution. The groundtruth phase offset between two beams (on two arrays) in the oracle solution is obtained using the method in ACO [41]. As we mentioned in Sec. 4, our radio only has a 2-bit phase shifter, which makes the measured impact of phase error at least the impact of 90 degrees phase difference. So we intentionally deviate the predicted phase from ground-truth by $\frac{\pi}{2}$, π and $\frac{3\pi}{2}$. Fig. 15 plots the normalized throughput gap generated different deviations. Our previous experiment already showed that X-Array can predict phase accurately with a median phase error of only about 0.5 in radius. This is much lower than the case with

phase error of $\pi/2$ which only causes about a median normalized throughput gap of 15%. Since $0.5 < \frac{\pi}{2}$, it is clear that X-Array suffers much less normalized throughput gap (around 12%) compared to the $\frac{\pi}{2}$ phase error case. In contrast, recall periodic probing has a phase prediction error between $\pi/2$ to π (Fig. 14), this translates into a significant median normalized throughput loss of 18% to 42%. This result will be further corroborated in the system level evaluation (Sec. 5.2). Note that higher throughput loss exhibits in outdoor scenario than indoor, because there are fewer reflectors in outdoor scenario which may cause less NLoS to compensate the throughput gap.

5.1.3 Recovering From link Outage.

We compare X-Array with the *neighbor scan only* (NS) approach to check its effectiveness and overhead in recovering from blockage. Our experiment investigates indoor scenarios where blockage often occurs due to human activities. We collect 10 trials with random client locations in each environment and evaluate the *normalized throughput gap* when blockage occurs. Fig. 16 plots the CDF across all the experiment trials. It is observed that X-Array experiences almost no throughput gap in 96% and 93% of the time in complex indoor and indoor open space, respectively. In contrast, NS suffers from throughput loss in over 60% of the cases, which can be up to 99.8% in complex indoor scenario. The indoor open space does not have as frequent sudden change as complex indoor scenario, thus more likely to find the strongest path in a shorter time, but the 80% normalized throughput gap is still large (about 99%).

5.2 System Level Evaluation

Overhead reduction. We compare the overall throughput of X-Array and periodic probing in open space and complex indoor/outdoor environment and conduct the experiments in 5 trials at different locations for each setting. The periodic probing scheme shares the same array/beam selection mechanism as X-Array, but aligns the phase based on different CSI feedback frequencies, *i.e.*, every 1 , $\frac{1}{2}$ and $\frac{1}{3}$ BI. Also the high speed results are obtained by predicting phase by X-Array and periodic probing using the method in Sec. 5.1.1 first, measuring the RSS of joint-panel beamforming with the predicted phases, and then mapping it to throughput. The box plot in Fig. 17 shows that X-Array has a near-zero median and 75-percentile throughput gap, and only several outliers in all the settings. In contrast, periodic probing suffers from around 20% to 57% normalized throughput gap. Interestingly, although high frequency periodic probing (*e.g.*, at $\frac{1}{3}$ -BI intervals) may achieve low normalized throughput gap at low moving speed, the normalized throughput gap increases dramatically when it comes to high-speed case. This is because the phase change rate may easily exceed the probing frequency under high mobility. The experiment verifies the importance of the predictive co-phasing of X-Array under different

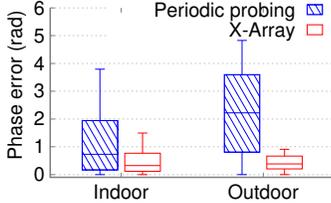


Figure 14: Phase prediction error.

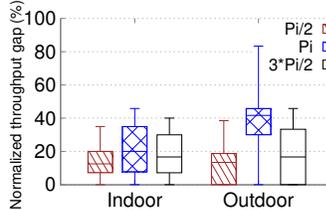


Figure 15: Impact of phase prediction error.

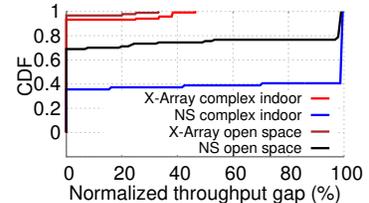


Figure 16: CDF of normalized throughput gap under blockage

moving speeds.

Coverage improvement. To verify whether X-Array effectively exploits the coverage advantages of APA (Sec. 2), we conduct experiments with different number of arrays and by disabling/enabling its design components in both typical indoor and outdoor scenarios. For indoor scenario, we place the client at certain locations and conduct measurements across the whole room. For outdoor scenario, we fix the AP on a lamppost (Fig. 10) and place the client at different distances, and repeat each distance setting in five different outdoor environments. We evaluate the cases with 2, 4, and 8 arrays running X-Array, and a case with 4 arrays but disabling the array selection and co-phasing (labeled as “4 arrays w/o switching”). According to the specification of our device, the TRP regulation constraint mandates that at most 2 arrays be turned on with full power. So when $N > 2$ arrays are active simultaneously, we reduce the transmit power per array to $2/N$. The result of indoor and outdoor scenarios are shown in Fig. 18 and Fig. 19, from which we can derive the following major insights: (i). Although 8-array achieves similar coverage as 4-array, the areas with high bit-rate links is much larger, thanks to arrays with partially overlapping FoVs providing co-phasing gains. Overall, with 8-arrays, X-Array eliminates all blind spots in the room, even for locations with thick walls, *i.e.*, it approximately achieves omni-directional coverage. (ii). In the indoor and some outdoor cases, the 4-array w/o switching performs even worse than 2-array running X-Array. This is due to two reasons. *First*, the former turns on all arrays, and wastes transmit power on arrays that may not provide any multipath diversity. Interestingly, we found that for 4-Array and 8-Array cases running X-Array, only 2 arrays are activated most of the time, implying it can intelligently select the best arrays rather than turning on all. *Second*, the phases between panels in 4-array w/o switching case are randomly chosen and not coherently aligned most of the time. Consequently, the throughput performance is harmed and becomes unstable, implying the importance of the co-phasing design.

6 DISCUSSION

APA represents a relatively new phased array architecture to establish high-performance mmWave networks. Our X-Array system has addressed several major challenges in APA, but many other design choices exist, which we discuss below.

Extension to multi-RF-chain mmWave MIMO. Emerging mmWave network standards such as 802.11ay and 5G NR support mmWave MIMO, *i.e.*, multiple RF chains each connecting to one phased array, sending multiple streams of data simultaneously to a single user (SU-MIMO multiplexing mode) or multiple users (MU-MIMO). As mentioned in Sec. 3.1, extension to multiple users in single-RF chain device can be realized using the built-in MAC protocol in 802.11ad like CSMA/TDMA based scheduling. Alternatively, mmWave MIMO can send the same stream of data across multiple RF chains to a single user to improve its SNR (SU-MIMO

diversity mode). Although APA has a single RF chain, it can be considered as a special case of SU-MIMO diversity. The AoD estimation, array/beam selection and outage handling mechanisms in X-Array can thus be directly applied to facilitate SU-MIMO. Its co-phasing implementation has been constrained by the 802.11ad hardware (2-bit phase resolution), yet the dynamic co-phasing formulation is general enough for future 802.11ay devices with phase weights. Even for the mmWave multi-user MIMO (MU-MIMO) case, X-Array’s three key mechanisms can serve as essential facilitating functions. Practical mmWave MU-MIMO protocols need to separate the analog beamforming and digital beamforming in two steps [45]. The former still heavily relies on the angular estimation and beam selection to narrow down the beam search space and then identify the best beam on each array for each user. These challenges are similar in nature to APA and can still benefit from the basic design components of X-Array. The detailed design that integrates X-Array in mmWave MIMO is left for our future work.

Dealing with more versatile client devices. Our current X-Array design assumes a simple single-array mobile client device. But future mmWave mobile devices may encompass APA to overcome users’ hand/body blockage. Accordingly, the array selection mechanism needs to be updated to coordinate the AP and client simultaneously. In addition, X-Array derives its decisions mainly based on the 802.11ad SLS, when the client turns to quasi-omni mode. Ideally, the client can further select its receiving beam after the AP’s array/beam are selected. This function is not implementable on our current 802.11ad device, but may be explored when an APA software radio becomes available.

Impact on higher layers Dynamic co-phasing of X-Array prevents most of the sudden drops of link quality (indicated by link RSS). Such variation will severely affect the effective throughput on the higher layer protocols and applications that are sensitive to instantaneous bandwidth estimation (*e.g.* TCP and DASH) [47, 48]. Such amplified impacts on higher layers are well studied and beyond the scope of our current work.

Multi-APA interference and spatial reuse. Our present work focuses on optimizing a single-cell mmWave network with a single AP. When multiple clients co-exist, the decision of each can be made independently, and the overhead will not increase in a noticeable way, because all of them share the SLS broadcast beacons from the AP. When multiple APs and clients coexist, X-Array may not sacrifice spatial reuse much since it tries to direct power towards the dominant AoA. Proper interference management schemes (*e.g.*, [49]) may still be needed, but are beyond the scope of this work.

7 RELATED WORK

In order to overcome the two main obstacles in realizing robust mmWave networking, *i.e.* mobility and blockage, recent research has explored efficient beam management algorithms, along with

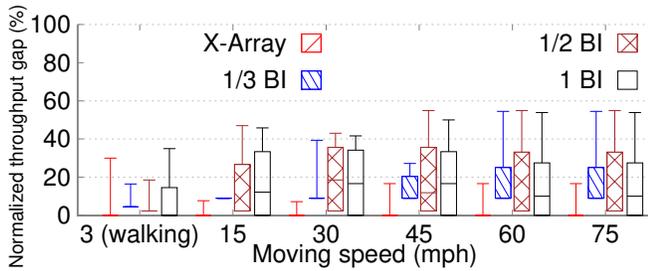


Figure 17: Normalized throughput gap due to co-phasing overhead.

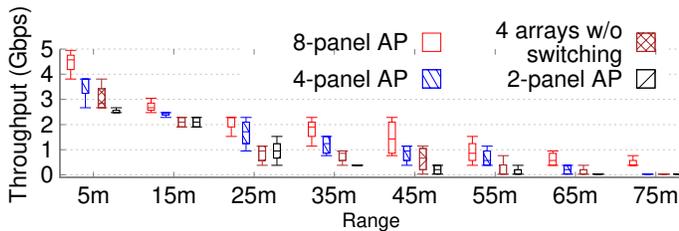


Figure 19: Multi-array outdoor range improvement.

new network architectures.

By reducing beam scanning overhead, the transmitter and receiver can quickly realign their beams, thus becoming resilient under channel dynamics. AgileLink [18] makes use of multi-arm beams and Hash function to identify the signal power along all spatial directions, and selects the beam along the strongest direction. BeamSpy [17] learns the correlation between beams offline, and prunes the beam search space to efficiently recover from blockage. UbiG [50] introduces an asymptotically efficient beam alignment algorithm that uses a few probeings to estimate the best beam. Most of such algorithms, along with many compressive sensing and statistical estimation algorithms [51–53], rely on the CSI as input, which requires non-trivial on-board memory space [33] and is unavailable on typical commodity mmWave radios. ACO [41] acquires CSI indirectly by measuring the RSS corresponding to different phase shift values of different groups antenna elements. Based on the CSI, it [41] estimates AoD according to [46] and form beams accordingly. However, the CSI acquisition process itself takes non-trivial overhead.

To overcome the FoV limitations of phased-arrays, Pia [24] leverages multiple cooperative APs, and switches to the appropriate AP whenever one is blocked. It uses motion and location sensors on mobile mmWave devices to overcome user mobility and orientation changes. Bouncenet [49] further addresses the spatial reuse when multiple APs and clients coexist. EMi [54] first reconstructs the reflection environment using mmWave sensing, and then intelligently places the APs to improve long-term network robustness

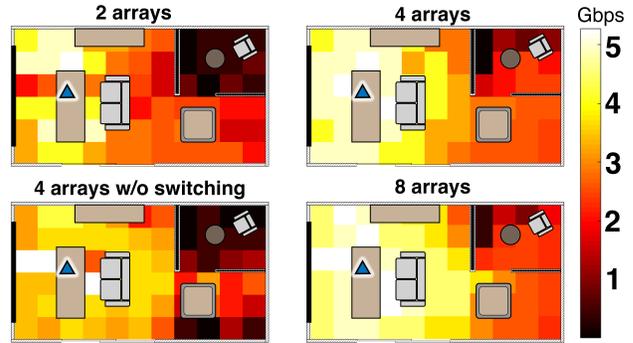


Figure 18: Coverage improvement in a room environment.

under random blockage and mobility. Beam-forecast and miDroid [19, 55] also leverages the environment information, and matches the measured CSI with ray-tracing simulated CSI in order to guide the beam selection. Yang *et al.* [56] leverages the mobile sensor data to adapt the beamwidth by choosing the beam from a multi-level codebook. Lister [57] makes use of luminaries information from lighting infrastructure to help with maintaining beam alignment and tracking mobility. Notably, none of the above work addresses the challenges related to the APA architecture (Sec. 2.2).

In emerging mmWave SU/MU-MIMO standards such as 802.11ay, the AP first performs legacy beam scanning and let clients report a set of potential beams for each phased array. Then the AP collects CSI feedback of the clients’ selected beams and further performs digital precoding to realize hybrid beamforming. Recent work designed algorithms to group the clients [58] or estimate the mmWave MIMO channel [59, 60], assuming detailed CSI feedback is available. As discussed in Sec. 6, X-Array addresses a different set of problems, although its design components can be transferable to single user mmWave MIMO.

8 CONCLUSION

We have explored APA as a new paradigm to simultaneously improve the mmWave coverage and link quality. Our X-Array solution framework builds on the 802.11ad standard and runs directly on commodity devices. Our experiments have verified X-Array’s advantages in terms of efficiency, coverage, and ability to rapidly recover from link outage. We believe X-Array marks an important step in making mmWave networks more resilient in dynamic and mobile scenarios.

ACKNOWLEDGEMENT

We sincerely thank our anonymous shepherd and the anonymous reviewers for their insightful feedback. This work was supported in part by the US National Science Foundation through NSF CNS-1854472, CNS-1506657, CNS-1518728, CNS-1617321, CNS-1925767, and the research funding from Ulsan National Institute of Science and Technology (UNIST) under Grant No. 1.180065.01.

REFERENCES

- [1] Federal Communications Commission, "FCC Promotes Higher Frequency Spectrum for Future Wireless Technology," <https://apps.fcc.gov>, 2015.
- [2] T. S. Rappaport, S. Sun, R. Mayzus, H. Zhao, Y. Azar, K. Wang, G. N. Wong, J. K. Schulz, M. Samimi, and F. Gutierrez, "Millimeter wave mobile communications for 5g cellular: It will work!" *IEEE access*, vol. 1, pp. 335–349, 2013.
- [3] The White House, "Unlocking the Promise of Broadband for All Americans." <https://www.whitehouse.gov/blog/2016/07/15/unlocking-promise-broadband-generate-gains-all-americans>, 2016.
- [4] Cisco Systems, Inc, "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2016-2021 White Paper," <https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/mobile-white-paper-c11-520862.html>, 2016.
- [5] O. Abari, D. Bharadia, A. Duffield, and D. Katabi, "Cutting the cord in virtual reality," in *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*. ACM, 2016, pp. 162–168.
- [6] K. Guan, D. He, A. Hrovat, B. Ai, Z. Zhong, and T. Kürner, "Challenges and chances for smart rail mobility at mmwave and thz bands from the channels viewpoint," in *2017 15th International Conference on ITS Telecommunications (ITST)*. IEEE, 2017, pp. 1–5.
- [7] S. E. Elayoubi, M. Fallgren, P. Spapis, G. Zimmermann, D. Martín-Sacristán, C. Yang, S. Jeux, P. Agyaopong, L. Campoy, Y. Qi, and S. Singh, "5G Service Requirements and Operational Use Cases: Analysis and METIS II Vision," in *European Conference on Networks and Communications (EuCNC)*, 2016.
- [8] M. Tercero, P. von Wrycza, A. Amah, J. Widmer, M. Fresia, V. Frascolla, J. Lorca, T. Svensson, M. Hamon, S. Destouet Roblot, A. Vijay, M. Peter, V. Sgardoni, M. Hunukumbure, J. Luo, and N. Vucic, "5G Systems: The mmMAGIC Project Perspective on Use Cases and Challenges Between 6–100 GHz," in *IEEE Wireless Communications and Networking Conference*, 2016.
- [9] C. Tranoris, S. Denazis, L. Guardalben, J. Pereira, and S. Sargento, "Enabling Cyber-Physical Systems for 5G Networking: A Case Study on the Automotive Vertical Domain," in *IEEE/ACM International Workshop on Software Engineering for Smart Cyber-Physical Systems (SESCPS)*, 2018.
- [10] M. Boers, B. Afshar, I. Vassiliou, S. Sarkar, S. T. Nicolson, E. Adabi, B. G. Perumana, T. Chalvatzis, S. Kavvadias, P. Sen, W. L. Chan, A. H. Yu, A. Parsa, M. Nariman, S. Yoon, A. G. Besoli, C. A. Kyriazidou, G. Zochios, J. A. Castaneda, T. Sowlati, M. Rofougaran, and A. Rofougaran, "A 16TX/16RX 60 GHz 802.11ad Chipset With Single Coaxial Interface and Polarization Diversity," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 12, 2014.
- [11] R. J. Mailloux, *Phased array antenna handbook*. Artech house, 2017.
- [12] P. Hannan, "The element-gain paradox for a phased-array antenna," *IEEE Transactions on Antennas and Propagation*, vol. 12, no. 4, pp. 423–433, 1964.
- [13] C. A. Balanis, *Antenna theory: analysis and design*. John Wiley & sons, 2016.
- [14] B. Li, Z. Zhou, W. Zou, X. Sun, and G. Du, "On the efficient beam-forming training for 60ghz wireless personal area networks," *IEEE Transactions on Wireless Communications*, vol. 12, no. 2, pp. 504–515, 2012.
- [15] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan, "60 ghz indoor networking through flexible beams: A link-level profiling," in *ACM SIGMETRICS Performance Evaluation Review*, vol. 43, no. 1. ACM, 2015, pp. 71–84.
- [16] A. Patra, L. Simić, and P. Mähönen, "Smart mm-wave beam steering algorithm for fast link re-establishment under node mobility in 60 ghz indoor wlans," in *Proceedings of the 13th ACM International Symposium on Mobility Management and Wireless Access*. ACM, 2015, pp. 53–62.
- [17] S. Sur, X. Zhang, P. Ramanathan, and R. Chandra, "BeamSpy: Enabling Robust 60 GHz Links Under Blockage," in *Proceedings of Usenix Conference on Networked Systems Design and Implementation (NSDI)*, 2016.
- [18] O. Abari, H. Hassanieh, M. Rodriguez, and D. Katabi, "Millimeter wave communications: From point-to-point links to agile network connections," ACM, pp. 169–175, 2016.
- [19] A. Zhou, X. Zhang, and H. Ma, "Beam-forecast: Facilitating mobile 60 ghz networks via model-driven beam steering," in *IEEE Conference on Computer Communications*. IEEE, 2017.
- [20] J. Palacios, D. De Donno, and J. Widmer, "Tracking mm-wave channel dynamics: Fast beam training strategies under mobility," in *IEEE Conference on Computer Communications*. IEEE, 2017, pp. 1–9.
- [21] T. Sowlati, S. Sarkar, B. G. Perumana, W. L. Chan, A. Papio Toda, B. Afshar, M. Boers, D. Shin, T. R. Mercer, W. Chen, A. Grau Besoli, S. Yoon, S. Kyriazidou, P. Yang, V. Aggarwal, N. Vakilian, D. Rozenblit, M. Kahrizi, J. Zhang, A. Wang, P. Sen, D. Murphy, A. Sajjadi, A. Mehrabani, E. Kornaros, K. Low, K. Kimura, V. Roussel, H. Xie, and V. Kodavati, "A 60-GHz 144-Element Phased-Array Transceiver for Backhaul Application," *IEEE Journal of Solid-State Circuits*, vol. 53, no. 12, 2018.
- [22] T. Sowlati, S. Sarkar, B. Perumana, W. L. Chan, B. Afshar, M. Boers, D. Shin, T. Mercer, W. Chen, A. P. Toda, A. G. Besoli, S. Yoon, S. Kyriazidou, P. Yang, V. Aggarwal, N. Vakilian, D. Rozenblit, M. Kahrizi, J. Zhang, A. Wang, P. Sen, D. Murphy, M. Mikhemar, A. Sajjadi, A. Mehrabani, B. Ibrahim, B. Pan, K. Juan, S. Xu, C. Guan, G. Geshvindman, K. Low, N. Kocaman, H. Eberhart, K. Kimura, I. Elgorriaga, V. Roussel, H. Xie, L. Shi, and V. Kodavati, "A 60GHz 144-element phased-array transceiver with 51dBm maximum EIRP and $\pm 60^\circ$ beam steering for backhaul application," in *IEEE International Solid - State Circuits Conference (ISSCC)*, 2018.
- [23] K. Ramachandran, N. Prasad, K. Hosoya, K. Maruhashi, and S. Rangarajan, "Adaptive beamforming for 60 ghz radios: Challenges and preliminary solutions," in *Proceedings of the 2010 ACM international workshop on mmWave communications: from circuits to networks*. ACM, 2010, pp. 33–38.
- [24] T. Wei and X. Zhang, "Pose information assisted 60 ghz networks: Towards seamless coverage and mobility support," in *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*. ACM, 2017, pp. 42–55.
- [25] Microtik, "60 GHz Access Point with 3 Titled Phased Arrays," https://mikrotik.com/product/wap_60gx3_ap/fndtn-specifications, 2017.
- [26] "Airfide inc," <http://airfidenet.com>, 2019.
- [27] W. Hong, K. Baek, Y. Lee, Y. Kim, and S. Ko, "Study and Prototyping of Practically Large-Scale mmWave Antenna Systems for 5G Cellular Devices," *IEEE Communications Magazine*, vol. 52, no. 9, 2014.
- [28] Y. Huang, Y. Li, H. Ren, J. Lu, and W. Zhang, "Multi-panel mimo in 5g," *IEEE Communications Magazine*, vol. 56, no. 3, pp. 56–61, 2018.
- [29] A. Ghosh, "5g new radio (nr): physical layer overview and performance," in *IEEE communication theory workshop*, 2018, pp. 1–38.
- [30] E. Onggosanusi, M. S. Rahman, L. Guo, Y. Kwak, H. Noh, Y. Kim, S. Faxer, M. Harrison, M. Frenne, S. Grant *et al.*, "Modular and High-Resolution Channel State Information and Beam Management for 5G New Radio," *IEEE Communications Magazine*, vol. 56, no. 3, 2018.
- [31] J. Liu, K. Au, A. Maaref, J. Luo, H. Baligh, H. Tong, A. Chassaingne, and J. Lorca, "Initial access, mobility, and user-centric multi-beam operation in 5g new radio," *IEEE Communications Magazine*, vol. 56, no. 3, pp. 35–41, 2018.
- [32] N. Song, P. Wen, H. Sun, and T. Yang, "Multi-panel based hybrid beamforming for multi-user massive mimo," in *GLOBECOM 2017-2017 IEEE Global Communications Conference*. IEEE, 2017, pp. 1–6.
- [33] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "Practical MU-MIMO User Selection on 802.11Ac Commodity Networks," in *Proceedings of ACM MobiCom*, 2016.
- [34] D. Carlson, "Breaking Through the Cost Barrier for Phased Arrays," *Microwave Journal*, 2018.
- [35] J. Zhang, X. Zhang, P. Kulkarni, and P. Ramanathan, "OpenMili: A 60 GHz Software Radio Platform With a Reconfigurable Phased-Array Antenna," in *ACM MobiCom*, 2016.
- [36] R. Zhao, T. Woodford, T. Wei, K. Qian, and X. Zhang, "M-Cube: A Millimeter-Wave Massive MIMO Software Radio," in *ACM MobiCom*, 2020.
- [37] X. Zhang, "Millimeter-Wave V2X Testbed," <http://m3.ucsd.edu/>, 2020.
- [38] IEEE Standards Association, "IEEE Standards 802.11ad-2012: Enhancements for Very High Throughput in the 60 GHz Band," 2012.
- [39] Federal Communications Commission, "FCC 13-112," <https://apps.fcc.gov/edoc-public/attachmatch/FCC-13-112A1.pdf>.
- [40] D. De Donno, J. P. Beltrán, D. Giustiniano, and J. Widmer, "Hybrid analog-digital beam training for mmwave systems with low-resolution rf phase shifters," in *2016 IEEE International Conference on Communications Workshops (ICC)*. IEEE, 2016, pp. 700–705.
- [41] J. Palacios, D. Steinmetzer, A. Loch, M. Hollick, and J. Widmer, "Adaptive codebook optimization for beam training on off-the-shelf ieee 802.11 ad devices," in *Proceedings of the ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2018.
- [42] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "WiFi-Assisted 60 GHz Wireless Networks," in *ACM MobiCom*, 2017.
- [43] D. Steinmetzer, D. Wegemer, and M. Hollick. (2017) Talon tools: The framework for practical ieee 802.11ad research. [Online]. Available: <https://seemoo.de/talon-tools>
- [44] IEEE 802.11ay Task Group, "Status of Project IEEE 802.11ay," http://www.ieee802.org/11/Reports/tgay_update.html, 2018.
- [45] P. Zhou, K. Cheng, X. Han, X. Fang, Y. Fang, R. He, Y. Long, and Y. Liu, "IEEE 802.11ay-based mmWave WLANs: Design Challenges and Solutions," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 3, 2018.
- [46] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE transactions on antennas and propagation*, vol. 34, no. 3, pp. 276–280, 1986.
- [47] X. Xie, X. Zhang, and S. Zhu, "Accelerating mobile web loading using cellular link information," in *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys '17. New York, NY, USA: ACM, 2017, pp. 427–439.
- [48] X. Xie, X. Zhang, S. Kumar, and L. E. Li, "PiStream: Physical Layer Informed Adaptive Video Streaming over LTE," in *Proc. of ACM MobiCom*, 2015.
- [49] S. Jog, J. Wang, J. Guan, T. Moon, H. Hassanieh, and R. R. Choudhury, "Many-to-many beam alignment in millimeter wave networks," in *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2019.
- [50] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "Towards scalable and ubiquitous millimeter-wave wireless networks," in *the 24th Annual International Conference on Mobile Computing and Networking*. ACM, 2018, pp. 257–271.
- [51] Z. Marzi, D. Ramasamy, and U. Madhow, "Compressive channel estimation and tracking for large arrays in mm-wave picocells," *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 3, pp. 514–527, 2016.
- [52] D. Ramasamy, S. Venkateswaran, and U. Madhow, "Compressive adaptation of large steerable arrays," in *ITA*, 2012, pp. 234–239.
- [53] —, "Compressive tracking with 1000-element arrays: A framework for multi-gbps mm wave cellular downlinks," in *2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2012, pp. 690–697.
- [54] T. Wei, A. Zhou, and X. Zhang, "Facilitating Robust 60 GHz Network Deployment

- by Sensing Ambient Reflectors,” in *Proc. of USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2017.
- [55] A. Zhou, S. Xu, S. Wang, J. Huang, S. Yang, T. Wei, X. Zhang, and H. Ma, “Robot Navigation in Radio Beam Space: Leveraging Robotic Intelligence for Seamless MmWave Network Coverage,” in *Proceedings of ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2019.
- [56] Z. Yang, P. H. Pathak, Y. Zeng, and P. Mohapatra, “Sensor-Assisted Codebook-Based Beamforming for Mobility Management in 60 ghz WLANs,” in *IEEE International Conference on Mobile Ad Hoc and Sensor Systems*, 2015.
- [57] M. K. Haider, Y. Ghasempour, D. Koutsonikolas, and E. W. Knightly, “Lis-teer: mmWave Beam Acquisition and Steering by Tracking Indicator LEDs on Wireless APs,” in *Proceedings of ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2018.
- [58] Y. Ghasempour and E. W. Knightly, “Decoupling beam steering and user selection for scaling multi-user 60 ghz wlans,” in *Proceedings of the 18th ACM International Symposium on Mobile Ad Hoc Networking and Computing*. ACM, 2017, p. 10.
- [59] Y. Ghasempour, M. K. Haider, C. Cordeiro, D. Koutsonikolas, and E. Knightly, “Multi-Stream Beam-Training for mmWave MIMO Networks,” in *Proceedings of ACM MobiCom*, 2018.
- [60] A. Alkhateeb, O. El Ayach, G. Leus, and R. W. Heath, “Channel Estimation and Hybrid Precoding for Millimeter Wave Cellular Systems,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 5, 2014.