FTP: Enabling Fast Beam-Training for Optimal mmWave Beamforming

Wei-Han Chen[§], Xin Liu[§], Kannan Srinivasan, and Srinivasan Parthasarathy
The Ohio State University
chen.7323@osu.edu, liu.10663@osu.edu, kannan@cse.ohio-state.edu, srini@cse.ohio-state.edu

IEEE INFOCOM 2024 - IEEE Conference on Computer Communications | 979-8-3503-8350-8/24/$31.00 ©2024 IEEE | DOI: 10.1109/INFOCOM52122.2024.10621112

***Abstract*—To maximize Signal-to-Noise Ratio (SNR), it is necessary to move beyond selecting beams from a codebook. While the state-of-the-art approaches can significantly improve SNR compared to codebook-based beam selection by exploiting the globally-optimal beam, they incur significant *beam-training overhead*, which limits the applicability to large-scale antenna arrays and the scalability for multiple users. In this paper, we propose *FTP*, a highly-scalable beam-training solution that can find the globally-optimal beam with minimal *beam-training overhead*. *FTP* works by estimating per-path direction along with its complex gain and synthesizes the globally-optimal beam from these parameters. Our design significantly reduces the search space for finding such path parameters, which enables *FTP* to scale to large-scale antenna arrays. We implemented and evaluated *FTP* on a mmWave experimental platform with 32 antenna elements. Our results demonstrate that *FTP* achieves optimal SNR performance comparable with the state-of-the-art while reducing the *beam-training overhead* by 3 orders of magnitude. Under simulated settings, we demonstrate that the gain of *FTP* can be even more significant for larger antenna arrays with up to 1024 elements.**

***Index Terms*—mmWave, beamforming, overhead**

I. Introduction

The availability of large bandwidth in mmWave frequency bands (e.g., 30-300GHz) is promising to meet the increasing demands for high data rates. This is particularly crucial in today's world where high data rate applications such as virtual and augmented reality, high-definition video streaming, and remote healthcare are becoming increasingly popular. To support these cutting-edge applications, new communication standards are also incorporating mmWave technologies into IEEE 802.11 wireless LAN [1] and 5G networks [2].

However, the high path loss in mmWave frequency bands poses significant challenges. In mmWave communications, beamforming is required to compensate for such high path loss. A transmitter (TX) needs to first align its beam with the receiver (RX) through a beam-training process before it can start data transmission. Although various beam-training methods [1], [3]–[12] have been proposed, these approaches cannot maximize the SNR because their beam selection is limited to a pre-defined codebook and therefore cannot find the globally-optimal beam corresponding to the channel. Such an optimal beam amplifies transmitted signals along the directions of the physical paths and ensures constructive interference at the RX. Thus, to maximize SNR, it is necessary to move beyond selecting beams from a codebook. State-of-the-art (*ACO* [13] and *UbiG* [14]) can significantly improve SNR (∼2x [13]) compared to codebook-based beam selection by exploiting the globally-optimal beam.

While *ACO* [13] and *UbiG* [14] are effective in finding the globally-optimal beam, they incur significant *beam-training overhead*, which limits the applicability of such ideas as we move to large-scale antenna arrays and across multiple users. The *beam-training overhead* consists of two parts: (1) *probing overhead* (time needed by the TX to send beam probes and collect feedback for decision-making) and (2) *computational overhead*. The *probing overhead* of *ACO* scales linearly with the antenna array size N ($\approx 5N$), which becomes intractable for large-scale antenna arrays. *UbiG*'s use of a genetic algorithm [15] can take a long time to converge[1], resulting in significant *computational overhead*. Although customized hardware can considerably speed up the convergence of a genetic algorithm, it is not available on Commercial-Off-the-Shelf (COTS) mmWave devices.

In this paper, we propose *FTP*, a highly-scalable beam-training solution that can find the globally-optimal beam with minimal *beam-training overhead*. *FTP* works by estimating per-path direction along with its complex gain and synthesizes the globally-optimal beam from these parameters. Our design significantly reduces the search space for finding such path parameters, which enables *FTP* to scale to large arrays.

For example, a path with an Angle of Departure (AoD) of θ and a complex gain of $ae^{j\phi}$ results in an observed channel of $h = ae^{j\phi}g(\theta)$, where $g(\theta)$ denotes the directional gain of the transmit beam. Therefore, the observed channel is a function of both the gain of the physical path and the directional gain of the employed beam. To estimate individual components θ and $ae^{j\phi}$ from the ensemble is non-trivial. A state-of-the-art algorithm [14] searches through all possible combinations of ($ae^{j\phi}$,θ) to find the maximum match with the measured channels, which results in significant *computational overhead*. Our **key insight** is that it is possible to solve for θ directly from the ensemble and $ae^{j\phi}$ becomes trivial to obtain once θ is known. We observe that as the TX changes the transmit beams, the gain of the underlying physical path $ae^{j\phi}$ remains the same while the directional gain $g(\theta)$ changes as a function of θ. Since $g(\cdot)$ is known to the TX, this means that we can estimate the direction θ from the variation patterns of the observed

[§]Both authors contributed equally to the paper

[1]Up to a few seconds on an Intel i7-7700 CPU.

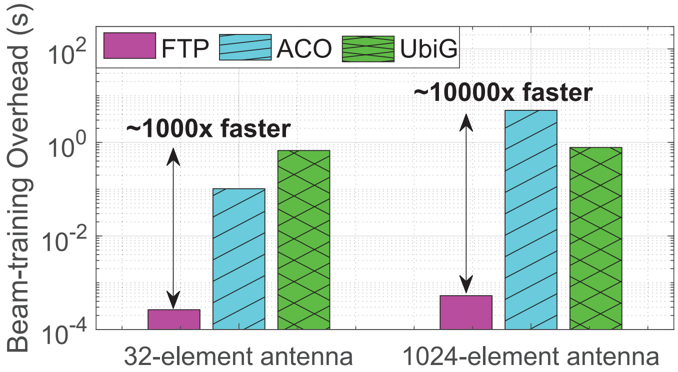

Fig. 1: Our experimental results show that *FTP* is 3 orders of magnitude faster than *UbiG* on a 32-element antenna array. Our simulation results show that *FTP* is 4 orders of magnitude faster than *ACO* on a 1024-element antenna array.

channels. By solving for the direction first, we can significantly reduce the *beam-training overhead* of finding the globally-optimal beam. In Fig. 1, our experimental results show that *FTP* is up to **3 orders of magnitude faster** [2] than *UbiG* on a 32-element antenna array (detailed in Sec. VI-E). Moreover, based on the insights and data from experiments on the smaller antenna array, we conduct simulations using a 1024-element antenna array to further demonstrate that the gain of *FTP* improves even further with larger antenna arrays. To realize *FTP*, however, several challenges need to be overcome.

First, *FTP* must efficiently and accurately estimate path directions from the channel variations. This step is crucial as it has a direct influence on the estimation accuracy of the complex gain for each path. By leveraging the sparsity of the mmWave channel [14], [16]–[18], we can formulate our path direction estimation as a compressive sensing (CS) problem. However, standard CS techniques [19]–[21] can not be applied here because the channel measurements across beam probes are phase-incoherent due to the Carrier Frequency Offset (CFO) on COTS mmWave devices [1], [2]. While a non-coherent CS design [22] addressed the phase-incoherence, it still has limitations to identify directions in multipath scenarios. Our **key observation** is that we can distinguish paths based on their Time of Flight (ToF) in the Channel Impulse Response (CIR) and individually solve for their direction, which enables us to break away from the limitations discussed above. By divide-and-conquer, we restrict the number of paths to 1 for each sub-problem and can thereby leverage the CS technique to efficiently and accurately estimate individual path direction.

Second, *FTP* must reliably distinguish paths in the CIR, which in turn ensures accurate path direction estimation. Two paths that have lengths within the distance resolution (inversely proportional to the system bandwidth) can not be distinguished in the CIR. Moreover, even the two paths are spaced more than the distance resolution, the observed channels can deviate from their true values due to the interaction between them. To understand this, the Nyquist sampling theorem [23] tells us that the digital representation of the CIR is the superposition of the two delayed and attenuated *sinc* functions centered at their ToF being measured at points spaced by the sampling interval. Therefore, two peaks that are farther away in the CIR affect each other less as the ripple of the *sinc* function slowly fades. To reliably distinguish paths in the CIR, we formulate it as an L1-norm regularized least-squares optimization problem, which is solved by convex optimization techniques [24], [25]. This optimization step results in improved accuracy and reliability of path direction identification.

Third, *FTP* addresses how to obtain coherent complex gain for each path such that we can guarantee constructive interference at the RX. Although channel measurements across beam probes are phase-incoherent due to the CFO on COTS mmWave devices [1], [2], we observe that the relative phases for each path within the same beam probe are still coherent. The reason is that the relative phases only depend on the difference between path lengths and they remain the same across beam probes. Furthermore, to ensure constructive interference at the RX, we only need the relative complex gains among the paths instead of their absolute values. Therefore, we measure the relative complex gains among paths in the same beam such that they are phase-coherent.

We note that *FTP* is complementary to beam-tracking solutions [26]–[29] that deal with link maintenance for blockage and mobility events (detailed in Sec. II). For example, mmReliable [28] can benefit from taking accurate initial input from *FTP* during beam-training and maintain optimal beam-tracking during data transmission.

We implemented and evaluated *FTP* on a mmWave experimental platform [30]. Our results demonstrate that *FTP* achieves optimal SNR performance comparable with state-of-the-art approaches [13], [14] while reducing the *beam-training overhead* by 3 orders of magnitude. Under simulated settings, we demonstrate that the gain of *FTP* can be even more significant for larger arrays with up to 1024 elements.

The contributions of this paper are summarized as follows:
(1) To the best of our knowledge, *FTP* is a mmWave beam-training approach that can find the globally-optimal beam with minimal overhead.
(2) We present the design of *FTP* that significantly reduces the search space for finding the globally-optimal beam.
(3) We implemented *FTP* on a mmWave experimental platform with 32 antenna elements and extensively evaluated it under realistic settings (Sec. VI). We demonstrate that *FTP* can scale to large antenna arrays.

II. Related work

FTP is a beam-training solution that enables a TX to efficiently align its beam with the RX in the beam-training phase prior to data transmission. During data transmission, beam-tracking techniques can be deployed to maintain the communication link for blockage and mobility events.

[2]While the absolute computation time is hardware-specific, the relative gain depends on the difference of algorithmic complexity and is hardware-independent.

■ **mmWave beam-training.** The IEEE 802.11ad standard [1] deploys a hierarchical beam search and it can take up to few seconds to find the correct beam. Recent compressive sensing solutions [20], [21] and AgileLink [18] enable a TX to align its beam with logarithmic number of measurements. However, they either require phase-coherence across beam probes [20], [21] or a customized phased array antenna [18] that are not available on COTS mmWave radios. Different from these approaches, *FTP* can work with COTS mmWave radios and do not require phase-coherence across measurements. Non-coherent compressive sensing desgins [7], [22], [31] were proposed and shown possible to be implemented on COTS mmWave radios. However, it still has limitations to align beams in multipath scenarios whereas *FTP* can find the optimal beam in such environments. Machine learning-based approaches [4], [6] aimed to train a deep neural network model to quickly select the best beam from a codebook. However, these methods can not go beyond selecting beams from a codebook whereas *FTP* can find the globally-optimal beam corresponding to the channel. In another line of work, researchers have considered using information from external sensors to assist beam selection, such as light sensors [32], cameras [9], MIMO WiFi radios [3], [10], and LiDAR [5]. Different from these methods, *FTP* identifies the optimal beam using in-band mmWave signals only and does not require extra sensors. When multiple APs and clients coexist, BounceNet [12] proposed an interference-aware beam-training algorithm that exploits spatial reuse to scale network throughput. Our work is complementary to such technique and can benefit from enabling more concurrent transmissions to further scale the gain of *FTP*.

ACO [13] and *UbiG* [14] are the state-of-the-art that can find the globally-optimal beam on COTS mmWave devices. *ACO* aims to measure the channel for each antenna element from the received power measurements and calculates the optimal beam from the full channel. However, the *probing overhead* for *ACO* scales linearly with antenna array size and it becomes intractable for large-scale antenna arrays. *UbiG* solves for parameters of each physical path through a genetic algorithm and assembles the optimal beam from the estimated parameters. However, *UbiG*'s algorithmic complexity incurs significant *computational overhead*, which limits its deployability on COTS mmWave devices. In contrast, *FTP* can find the globally-optimal beam while being orders of magnitude faster than *ACO* and *UbiG*, which enables *FTP* to be applied to large-scale antenna arrays and scale for multiple users.

■ **mmWave beam-tracking.** mmChoir [26] proactively mitigated blockage by joint transmissions from multiple APs to the client. MOCA [33] proactively identifies and adapts to link impairments due to user mobility through estimating link quality for selected beams prior to each data transmission. X-Array [27] leverages a multi-array architecture to efficiently find alternative arrays/beams when blockage occurs. Last, mmReliable [28] achieves link maintenance by tracking translational and rotational motions of the client. Complementary to these beam-tracking techniques that adpat beams during data transmission, *FTP* is a beam-training solution that enables a TX to efficiently find the globally-optimal beam during beam-training phase.

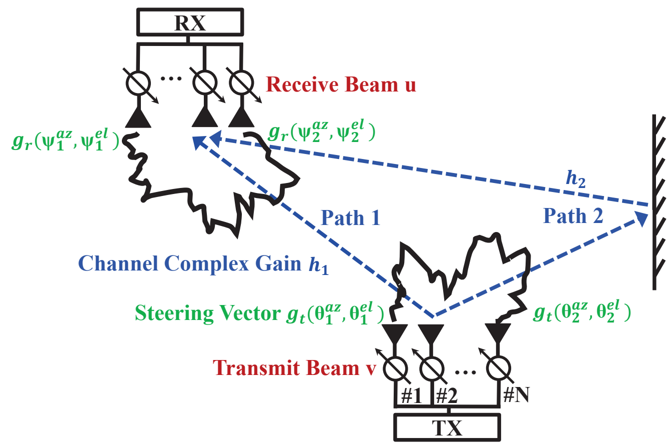

Fig. 2: The mmWave signal propagation from the TX to the RX involves various elements, including transmit/receive beams, steering vectors, and channel complex gains.

■ **Hybrid beamforming.** For Multiple-Input Multiple-Output (MIMO) mmWave systems, many techniques have been proposed [34]–[36] to jointly optimize signals at each phased array antenna for digital beamforming. However, they require multiple phased arrays whereas *FTP* focuses on analog beamforming systems with a single Radio Frequency (RF) chain. Our work is complementary to hybrid beamforming and can benefit from having multiple RF chains to align many optimal beams for multiple users.

III. Background on mmWave Channel

Obtaining mmWave channel information is crucial for deriving the optimal beam. Therefore, it is necessary to conduct a thorough analysis of the channel.

(1) *mmWave Communication.* We first present a mmWave communication model that enables us to analyze the relationship between the transmit beam at the TX and the received signal at the RX. Fig. 2 depicts the process of mmWave communication, which involves a few distinct steps: (i) TX Beamforming. The TX is equipped with a phased array antenna consisting of N elements. It applies various complex gains, including phase shifts and amplitudes, to each antenna element to form a beam in the desired direction. The combination of these complex gains is referred to as a beam. (ii) Signal Propagation. The transmitted signals travel through the wireless channel and may reflect off of an obstacle to reach the RX. (iii) RX Receiving. The RX is also equipped with a phased array antenna. A receive beam is applied to the antenna to boost the received signal. The whole process is given by:

$$\mathbf{y} = \mathbf{u}^H \mathbf{H} \mathbf{v} \mathbf{x} + \mathbf{u}^H \mathbf{n} \tag{1}$$

where $\mathbf{v}$ and $\mathbf{u}$ represent the transmit beam and receive beam, respectively. $\mathbf{y}$, $\mathbf{x}$, and $\mathbf{n}$ are the received signal at the RX, the transmitted signal from the TX, and the noise. $\mathbf{H}$ denotes the mmWave channel. Based on Eqn. 1, the CIR is given by:

$$\mathbf{p} = \mathbf{u}^H \mathbf{H} \mathbf{v} \tag{2}$$

Our design objective is to maximize the signal strength by optimizing $\mathbf{v}$ to the optimal beam. The received signal strength can be expressed as $|\mathbf{p}|^2$. By applying the Cauchy-Schwarz inequality [37], the optimal beam is given by:

$$\mathbf{v}^* = \mathbf{H}^H \mathbf{u} / ||\mathbf{H}^H \mathbf{u}|| \tag{3}$$

where division of the norm keeps the power constraint.

Thus, to obtain the optimal beam $\mathbf{v}^*$, all we need to know is $\mathbf{u}^H \mathbf{H}$, which is the *transmitting channel* viewed by the TX.

(2) *mmWave Transmitting Channel.* To fully exploit the potential of an optimal beam, we focus on the multipath scenarios in practical mmWave communications. Specifically, as shown in Fig. 2, the *transmitting channel* can be represented by a set of elements along paths between the TX and RX:

$$\mathbf{u}^H \mathbf{H} = \mathbf{u}^H \sum_{k=1}^{K} h_k g_r(\psi_k^{az}, \psi_k^{el}) g_t^H(\theta_k^{az}, \theta_k^{el}) \tag{4}$$

where K represents the number of paths. h_k is the channel complex gain along the k-th path. For the k-th path, the Angle of Arrival (AoA) (including the azimuth and elevation angles) is given by $(\psi_k^{az}, \psi_k^{el})$, while the Angle of Departure (AoD) is given by $(\theta_k^{az}, \theta_k^{el})$. $g_r(\cdot)$ and $g_t(\cdot)$ represent the steering vector functions at the TX and RX, respectively.

We note that the RX uses a fixed beam $\mathbf{u}$ during TX beam training, resulting in $\mathbf{u}$ and $g_r(\psi_k^{az}, \psi_k^{el})$ remaining constant across CIRs. h_k is also fixed for a specific path. Since the product of $\mathbf{u}^H h_k g_r(\psi_k^{az}, \psi_k^{el})$ is a constant complex value for a specific path, we can simplify Eqn. 4 as follows:

$$\mathbf{u}^H \mathbf{H} = \sum_{k=1}^{K} a_k e^{j\phi_k} g_t^H(\theta_k^{az}, \theta_k^{el}) \tag{5}$$

where $a_k e^{j\phi_k} = \mathbf{u}^H h_k g_r(\psi_k^{az}, \psi_k^{el})$ represents the total complex gain of the k-th path, which is the product of the receive beam, the channel complex gain, and the RX steering vector along that path.

Eqn. 1 and 5 reveal that the mmWave channel is a combination of various elements. Exploring these different elements can lead to diverse approaches for obtaining channel information, resulting in varying *beam-training overhead*.

IV. *FTP*: Key Design Insight

In this section we present the key facet underpinning *FTP*'s design. We begin by noting that prior research has approached the problem of channel measurement in a variety of ways. We systematically discuss some of these prior approaches before discussing the key insight that *FTP* leverages.

Antenna Perspective: First, from the perspective of antenna, the mmWave channel is an assembly of the distinct channels of each antenna element. We can express this concept by rewriting Eqn. 1:

$$\mathbf{y} = \sum_{n=1}^{N} [0 \cdots, (\mathbf{u}^H \mathbf{H})_n, \cdots 0][0 \cdots, \mathbf{v}_n, \cdots 0]^T \mathbf{x} + \mathbf{u}^H \mathbf{n} \tag{6}$$

where $(\mathbf{u}^H \mathbf{H})_n$ is the n-th element of *transmitting channel* $\mathbf{u}^H \mathbf{H}$, while $\mathbf{v}_n$ is the complex gain of the n-th antenna element. We note that both $\mathbf{u}^H \mathbf{H}$ and $\mathbf{v}$ in Eqn. 1 are complex vectors with sizes $1 \times N$ and $N \times 1$, respectively.

Eqn. 6 reveals that it is feasible to measure $(\mathbf{u}^H \mathbf{H})_n$ using $[0 \cdots, \mathbf{v}_n, \cdots 0]^T$ as the probing beam, as validated by *ACO* [13]. As a result, to measure the channel corresponding to each antenna element, this approach results in a *probing overhead* that scales linearly with N ($\mathcal{O}(N)$), leading to significant *probing overhead* for large-scale antenna arrays.

Path Perspective: Second, an alternative approach relies on a path-based approach. Eqn. 5 reveals that the mmWave channel is a combination of each individual path's channel, which in turn is jointly determined by 4 unknowns (a_k, ϕ_k, θ_k^{az}, θ_k^{el}). Since the four unknowns are considered to be tightly coupled, *UbiG* [14] extracts them by solving the following optimization problem:

$$\{a_k, \phi_k, \theta_k^{az}, \theta_k^{el}\}^* = \underset{\{a_k, \phi_k, \theta_k^{az}, \theta_k^{el}\}}{\operatorname{argmin}} ||\{(\mathbf{u}^H \mathbf{H})_k \mathbf{v}_m - \mathbf{p}_{(m,k)}\}_{m=1}^{M}||^2 \tag{7}$$

where $(\mathbf{u}^H \mathbf{H})_k$ represents the *transmitting channel* along the k-th path. $\mathbf{v}_m$ denotes the m-th probing beam, and the number of probing beam is M. $\mathbf{p}_{(m,k)}$ is the measured CIR of the k-th path from the m-th probing beam.

However, Eqn. 7 is a non-convex non-linear optimization problem. *UbiG* solves it by using a genetic algorithm with orthogonal matching pursuit. The large search space for finding the global optimum of this problem leads to a significant *computational overhead*, which can take a long time to converge (up to a few seconds on an Intel i7-7700 CPU).

■ ***FTP*'s Key Insight.** We observe that among the four unknowns, (θ_k^{az}, θ_k^{el}) can be solved separately from $c_k = a_k e^{j\phi_k}$. To better understand our observation, we can rewrite Eqn. 7 as follows:

$$\begin{aligned}(\theta_k^{az}, \theta_k^{el})^* &= \underset{(\theta_k^{az}, \theta_k^{el})}{\operatorname{argmin}} ||\{c_k g_t^H(\theta_k^{az}, \theta_k^{el}) \mathbf{v}_m - \mathbf{p}_{(m,k)}\}_{m=1}^{M}||^2 \\ &= \underset{(\theta_k^{az}, \theta_k^{el})}{\operatorname{argmin}} ||\{g_t^H(\theta_k^{az}, \theta_k^{el}) \mathbf{v}_m - \mathbf{p}_{(m,k)}\}_{m=1}^{M}||^2\end{aligned} \tag{8}$$

Since $a_k e^{j\phi_k}$ is a constant value for k-th path and does not depend on (θ_k^{az}, θ_k^{el}), it can be considered as a scaling factor in Eqn. 8 and be ignored during the optimization process without affecting the final results. Therefore, we can simplify the optimization problem by removing this constant factor from the objective function and focus solely on finding the optimal values of (θ_k^{az}, θ_k^{el}). Moreover, once (θ_k^{az}, θ_k^{el}) are determined, $a_k e^{j\phi_k}$ can be trivially derived. This approach significantly reduces the search space for the global optimum compared to Eqn. 7.

V. *FTP* System Design

Based on *FTP*'s key insight, the optimal beam can be obtained in two steps: (1) determining the path directions and (2) deriving the optimal beam. To accomplish this, *FTP* employs the compressive sensing (CS) technique to estimate the path directions. However, this is challenging due to the limitations of existing CS solutions on COTS mmWave devices, which can only identify the direction of the dominant path but not the directions of multipath. In this section, we first describe how to perform CS-based spatial probing and detail how we can leverage the CS technique to identify the directions of multipath. Finally, we explain how to derive the optimal beam based on the identified path directions.

A. Perform Spatial Probing

FTP utilizes CS-based spatial probing to measure the mmWave channel. The basic idea is that the TX utilizes pseudorandom beams to send preamble signals in different directions. Preamble signals are pre-defined signals that are known to both TX and RX, and they enable the CIR estimation with high accuracy. The reason why pseudorandom beams are utilized is two-fold: (1) they can sample the spatial channel effectively since they are uncorrelated to each other. (2) Such beams do not require specialized hardware to generate and only require coarse phase and amplitude control for each antenna element. This makes it feasible to be implemented on COTS mmWave devices with low-cost hardware [7], [38].

The m-th probing beam is set as:

$$\mathbf{v}_m = [1, e^{j\theta_1^m}, e^{j\theta_2^m}, ..., e^{j\theta_{N-1}^m}] \tag{9}$$

where $[\theta_1^m, \theta_2^m, \dots, \theta_{N-1}^m]$ represent the phase shifts, which are independent and identically distributed random variables from a uniform distribution on $\{0, \pi/2, \pi, 3\pi/2\}$. Fig. 3a shows an example of the generated probing beams.

■ **Probing Overhead.** By utilizing CS-based spatial probing, the *probing overhead* of *FTP* is bounded by $\mathcal{O}(K \log N)$, which scales logarithmically with the number of antenna elements (N) and linearly with the number of paths (K). Since the mmWave channel is sparse, typically $K \leq 3$ [14], [16]–[18], this means that *FTP* can easily scale to large antenna arrays.

■ **Feedback Overhead.** After the RX receives the probing beams, it calculates the correlation value between the received signal and the transmitted preamble. This correlation value is a measure of the CIR, where paths will show as peaks, as shown in Fig. 3b. The RX then sends the correlation peaks back to the TX as feedback. Current standards like IEEE 802.11ad have feedback mechanisms built into their beam-training protocols. *FTP* can collect the correlation peaks of the $K \log N$ beams by piggybacking them on one feedback packet, which takes less than 1 μs in IEEE 802.11ad.

B. Determine Path Directions

FTP first distinguishes paths based on their ToF in the CIR. Then, it individually applies the CS technique to each path to estimate the direction. This allows *FTP* to overcome the limitations of existing CS solutions on COTS mmWave devices in multipath scenarios.

Upon receiving the feedback from the RX, the TX needs to map the correlation peaks to each path based on their ToF, as shown in Fig. 3c. We first sort the peaks according to their ToF, and then group the peaks that belong to the same path. Specifically, for two peaks to be mapped to the same path, we define a threshold for the maximum time difference between them. Peaks that have a time difference greater than the threshold are mapped to different paths. Through this process, we can obtain the CIR samples for each path.

Once the CIR samples for each path are obtained, we individually apply the CS algorithm to each path to estimate its direction. Without loss of generosity, let's assume the TX is equipped with an Uniform Rectangular Array (URA). The TX steering vector is a known function that can be derived from the antenna array geometry. For the antenna element in the r-th row and c-th column, its contribution to the TX steering vector is a 2D sinusoid given by:

$$g_t^H(\theta_k^{az}, \theta_k^{el}) = \{e^{\frac{-2\pi d[cos(\theta_k^{az})sin(\theta_k^{el})r + sin(\theta_k^{az})sin(\theta_k^{el})c]}{\lambda}}\}_{r=1,c=1}^{R,C} \tag{10}$$

where d and λ represent the antenna spacing and the wavelength of the central frequency, respectively. R and C are the number of antenna elements in each row and column, respectively, and $N = R \times C$.

We leverage the CS algorithm to solve the optimization problem in Eqn. 8. The direction of the k-th path can be estimated by maximizing the following cost function:

$$\underset{(\theta_k^{az}, \theta_k^{el})}{\text{argmax}} \langle \frac{\{|g_t^H(\theta_k^{az}, \theta_k^{el})\mathbf{v}_m|\}_{m=1}^M}{||\{g_t^H(\theta_k^{az}, \theta_k^{el})\mathbf{v}_m\}_{m=1}^M||}, \frac{\{|\mathbf{p}_{(m,k)}|\}_{m=1}^M}{||\{\mathbf{p}_{(m,k)}\}_{m=1}^M||} \rangle \tag{11}$$

where $\langle \cdot \rangle$ is the inner product operation. $\mathbf{v}_m$ represents the m-th probing beam, and $M = K \log N$ denotes the total number of probes. $\mathbf{p}_{(m,k)}$ is the correlation peak of the k-th path obtained from the m-th probing beam.

It is noteworthy that *FTP* does not require phase-coherence across probes throughout the process of estimating the path directions as it only uses the magnitude part of the correlation peaks, i.e., $|\mathbf{p}_{(m,k)}|$. This makes *FTP* widely applicable to COTS mmWave devices.

■ **Extend *FTP* to mmWave 5G.** The key challenge of extending *FTP* to mmWave 5G lies in the reduced ability of distinguishing between different paths. Compared to the IEEE 802.11ad [1] that employs almost 2 GHz of RF bandwidth, mmWave 5G occupies significantly less RF bandwidth (up to 400 MHz [2]), resulting in reduced distance resolution. Thus, when the difference between the path lengths is less than 0.75 m, the correlation peaks in the time domain will overlap. Fig. 3b illustrates the problem of overlapping correlation values from different paths. If *FTP* were to rely solely on the

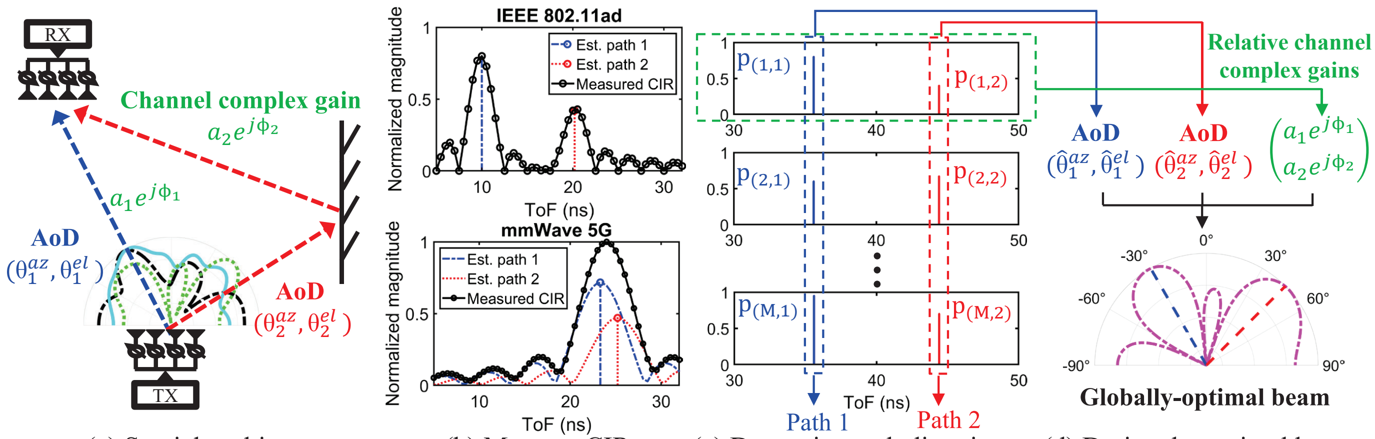

(a) Spatial probing (b) Measure CIR (c) Determine path directions (d) Derive the optimal beam

Fig. 3: *FTP*'s system overview. (a) First, CS-based spatial probing is utilized to sample the multipath channel. (b) Paths are detected as peaks in the CIR. Even when they overlap, we can distinguish them through an optimization process. (c) Path directions are estimated by individually performing CS on each path in the CIR. (d) Once directions are obtained, we compute the relative complex gains among paths from a particular beam that covers all paths. Then, *FTP* can construct the optimal beam from the estimated parameters.

method mentioned previously, it would be unable to accurately distinguish the direction of each path.

To address this challenge, we formulate it as an optimization problem and solve it by convex optimization. Based on the Nyquist sampling theorem [23], for a band-limited system, the time-domain signals can be perfectly reconstructed through *sinc* interpolation. Thus, the digital representation of the time-domain channel is the superposition of individual delayed and attenuated *sinc* centered at its ToF τ_k being sampled at points spaced by the sampling interval, which is given by:

$$h[n] = \sum_{k=1}^{K} c_k sinc(B(nT - \tau_k)) \tag{12}$$

where c_k, B, T, and τ_k denote the complex gain, the system bandwidth, the sampling interval, and the ToF for k-th path. Based on Eqn. 12, to recover (c_k, τ_k) from paths that overlap in the CIR, we can solve a matrix inverse problem by fitting *sinc* functions over the measured CIR. However, this results in an under-determined system of equations with infinitely many solutions. To overcome this, we leverage the sparsity of the mmWave channel and obtain the unique solution by solving the following L1-norm regularized least-squares optimization problem:

$$\mathbf{x}^* = \arg\min_{\mathbf{x}} ||\mathbf{D}\mathbf{x} - \mathbf{h}||_2^2 + \lambda||\mathbf{x}||_1 \tag{13}$$

where the matrix $\mathbf{D}$ consists of *sinc* functions centered at all ToF within the search range and the non-zero entries of $\mathbf{x}$ give (c_k, τ_k) for each path. Eqn. 13 can be solved efficiently by convex optimization techniques [24], [25].

C. Derive The Optimal Beam

After obtaining the path directions, *FTP* computes the complex gains $a_k e^{j\phi_k}$ of each paths and then derives the optimal beam $\mathbf{v}^*$.

However, the complex gains of each path are not phase-coherent across different CIRs. To address this issue, *FTP* leverages the phase-coherence within the same CIR to compute the relative complex gains between paths. Therefore, it is necessary for *FTP* to select a CIR that captures all paths. Fortunately, due to the sparsity of mmWave channels, finding such a CIR is often practical in real-world scenarios. We also consider the worst-case scenario where there is no CIR that can capture all paths. In such case, *FTP* amplifies the RF energy in the identified path directions and sends an additional probing beam specifically to capture the CIR from these paths. This guarantees *FTP* to have a CIR that captures all paths. Furthermore, this approach only requires sending an additional probing beam, resulting in negligible extra *probing overhead*.

Apart from the phase-incoherence issue, the CFO within a CIR needs to be removed as it introduces phase distortions. Fortunately, CFO removal is a standard operation done at the RX, which involves estimating the CFO from the preamble and applying phase shifts to the samples in the CIR. Therefore, the CIRs reported by the RX already have the CFO compensated.

Finally, the complex gain along k-th path is given by:

$$a_k e^{j\phi_k} = \frac{\mathbf{p}_{(z,k)}}{g_t^H((\theta_k^{az}, \theta_k^{el})^*)\mathbf{v}_z} \tag{14}$$

where $\mathbf{v}_z$ denotes the z-th probing beam that provides the selected CIR, and $\mathbf{p}_{(z,k)}$ denotes the selected CIR of the k-th path. Upon obtaining the complex gains, deriving the optimal beam becomes straightforward. Based on Eqn. 3 and 5, the optimal beam can be expressed as follows:

$$\mathbf{v}^* = \sum_{k=1}^{K} (a_k e^{j\phi_k})^H g_t((\theta_k^{az}, \theta_k^{el})^*) \tag{15}$$

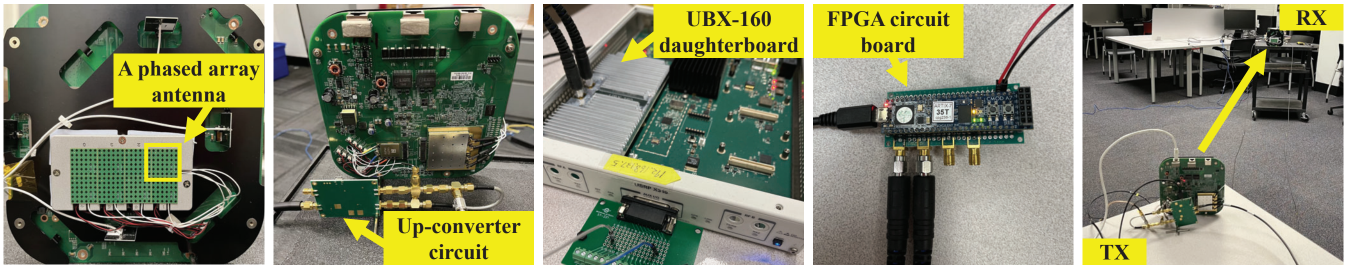

(a) Phased Array (b) 60 GHz Frequency (c) Baseband Processor (d) Beam Control (e) Indoor Office

Fig. 4: Implementation

VI. Evaluation

In this section, we provide a detailed description of our implementation, experimental setup, and performance evaluation. In addition, to prove the efficacy of our proposed beam-training approach, its results are compared with those from state-of-the-art techniques.

A. Implementation

To conduct a comprehensive evaluation, we implement *FTP*, 802.11ad [1], *ACO*, and *UbiG*.

■ ***FTP*.** To demonstrate the feasibility of *FTP* in practical scenarios, we implement it on the M-Cube platform [30], which utilizes a COTS phased array antenna from Airfide Inc [38], as shown in Fig. 4.

(1) *Phased Array. FTP* is equipped with a COTS phased array antenna with 32 steerable antenna elements, as shown in Fig. 4a. Each antenna element is controlled by a 2-bit phase-shifter, where four possible phase-shift values can be applied: $\{0,\ \pi/2,\ \pi,\ 3\pi/2\}$, and a one-bit amplitude control, i.e., an element can either be on or off.

(2) *60 GHz Frequency.* The phased array antenna used by *FTP* operates at 60 GHz for mmWave communications. To generate the necessary signals, the baseband signals from the baseband processor unit are first up-converted to 15 GHz intermediate frequency (IF) signals by an up-converter circuit board, as shown in Fig. 4b. The IF signals are then up-converted to 60 GHz by the antenna module.

(3) *Baseband Processing.* The M-Cube platform provides seamless integration with both USRP and MATLAB, enabling us to easily process baseband signals. Specifically, we use a USRP X310 radio with UBX-160 as our baseband processor unit, as shown in Fig. 4c, and perform signal processing using MATLAB. Due to the analog bandwidth limitation of the UBX-160 daughterboard, the bandwidth in our experiments is capped at 160 MHz.

(4) *Beam Control.* The phased array antenna is controlled in real-time for beam-switching by an FPGA, which sends beam-switching commands at specified timestamps as shown in Fig. 4d. To ensure accurate beam-switching, the control commands and the IF data signals are synchronized in time using the Automatic Transmit/Receive (ATR) function supported by the USRP. Moreover, customized beam patterns can be loaded into the device's internal memory via an Ethernet connection to modify the beam patterns.

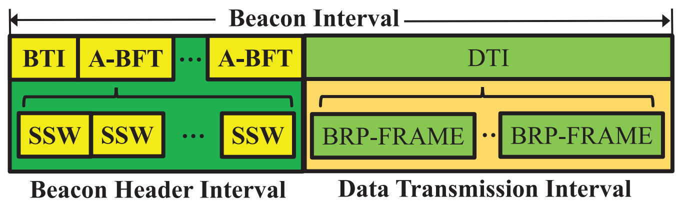

Fig. 5: The beam-training procedure specified by the IEEE 802.11ad standard.

■ **IEEE 802.11ad.** We also implement the IEEE 802.11ad on the M-Cube platform as a baseline beam-training approach, which is widely used for 60 GHz WLAN. The codebook we use is pre-configured by the device vendor [38] and contains 128 beams. This codebook size is in compliance with the IEEE 802.11ad standard [27]. To determine the optimal beam for the baseline, we probe through all available beams.

■ ***ACO*.** We also implement *ACO* on the M-Cube platform as the state-of-the-art beam-training approach, which can find the globally-optimal beam corresponding to the channel and maximizes the SNR. We select *ACO* as the benchmark for optimal SNR performance among state-of-the-art approaches.

■ ***UbiG*.** We also include *UbiG*, another state-of-the-art beam-training approach, in our experiments to compare the *FTP*'s *beam-training overhead* with that of *UbiG*.

B. Experiment Setup

■ **PHY & MAC settings.** We conduct experiments under different PHY and MAC settings of IEEE 802.11ad. It is worth noting that we do not require any hardware changes to the PHY layer or any modifications to the MAC layer, which demonstrates the potential integration of *FTP* with the existing standards.

(1) *PHY Layer.* We follow the PHY parameters specified in the IEEE 802.11ad standard to evaluate SNR and throughput. Our modulation coding scheme (MCS) range from 1 to 12, utilizing various types of modulation (BPSK, QPSK, and 16QAM), with each signal shifted by $\pi/2$ to avoid zero crossings in the I/Q domain. We employ low-density parity check (LDPC) codes for signal encoding, with code rates of 1/2, 5/8, 3/4, and 13/16 used depending on the data rate. Given

the 160 MHz bandwidth in our implementation, the achievable throughput can be up to 420 Mbps.

(2) *MAC Layer.* We adopt the IEEE 802.11ad MAC protocol to evaluate the *beam-training overhead* of *FTP* in terms of time. Specifically, the channel access is divided into periodic beacon intervals (BI), which last for 100 ms and consist of beacon header intervals (BHI) and data transmission intervals (DTI), as shown in Fig. 5. We perform *FTP* only during the BHI phase, which is divided into a beacon transmission interval (BTI) and up to 8 association beamforming training (A-BFT) slots. While beam refinement phase (BRP) frames can be used during the DTI phase to refine beam performance, we do not utilize them due to their limited deployment in COTS mmWave devices [7]. To demonstrate the baseline performance, we assume no collisions during the beam-training. However, *FTP* has a significantly lower overhead compared to state-of-the-art methods, which can lead to even better performance in the presence of collisions.

■ **Experimental Environment.** We conducted our experiments in an indoor office environment measuring 24×13 meters, which contains a variety of reflective objects such as metal cabinets, whiteboards, and concrete walls. The TX and RX were randomly placed at various locations in the office to cover both line-of-sight (LOS) and non-line-of-sight (NLOS) scenarios, which are typical for IEEE 802.11ad communications. Fig. 4e shows an example of placing the TX and RX in our experiments. At each location, the RX was fixed to receive with a fixed RX beam, while the TX iterated through different beam-training methods.

C. LOS Performance

In this section, we investigate the SNR and throughput performance of the three beam-training approaches with the same IEEE 802.11ad PHY settings in LOS scenarios. Since the TX uses an antenna array with 32-steerable elements, the number of probes required for *FTP* and *ACO* are 15 ($K \log N$) and 156 ($5N-4$), respectively. The 802.11ad, as the baseline approach, scans through the available 128 beams and selects the best performing one.

The resulting SNR performance is displayed in Fig.6a. The median SNR for *FTP*, *ACO*, and the baseline are 21.41, 21.42, and 17.65 dB, respectively. While *FTP*'s performance closely follows *ACO*, it requires significantly fewer probes. In contrast, the baseline has sub-optimal performance with almost 4 dB lower median SNR than *FTP* and *ACO*.

These results highlight the importance of generating constructive multi-beams in indoor environments where many reflectors, such as metal cabinets, whiteboards, and concrete walls, create strong alternative paths in addition to the LOS. To achieve optimal performance, it is necessary to synthesize the globally-optimal beam from per-path direction and complex gain parameters. *FTP* estimates these parameters and adapts its beams optimally to the channel, whereas the baseline approach can only select one of the beams from the predefined codebook and cannot fully leverage the opportunities provided by such multipath.

Moving on to the throughput performance shown in Fig. 6b, although all three beam-training approaches achieve similar median throughput of around 400 Mbps in LOS scenarios, *FTP* and *ACO* can attain maximum throughput performance twice as often as the baseline (80% versus 40%) due to their ability to consistently achieve maximum SNR.

D. NLOS Performance

In this section, using the same IEEE 802.11ad PHY settings, we investigate the SNR and throughput performance of *FTP*, *ACO*, and the baseline in NLOS scenarios. Again, the TX scans 15, 156, and 128 probes for *FTP*, *ACO*, and the baseline, respectively. The resultant SNR performance is shown in Fig. 6c. The median SNR for *FTP*, *ACO*, and the baseline are 13.58, 14.27, and 9.58 dB, respectively. *FTP*'s performance still closely follows *ACO* in NLOS scenarios with a median SNR loss of less than 1 dB. As expected, the baseline has sub-optimal performance with 4 dB lower median SNR than *FTP*. In NLOS scenarios, the transmitted signals can reach the receiver through a path by either reflecting off of a reflector or penetrating through an obstacle. Since the received signals are weaker in this case compared to LOS scenarios, it's even more important to leverage the multipath propagations and construct the corresponding multi-beam to maximize the SNR. The experiment results verify that *FTP* can accurately estimate per-path direction and relative complex gains among paths in NLOS settings and achieve optimal SNR performance. Next, Fig. 6d shows the throughput performance. The median throughput for *FTP*, *ACO*, and the baseline are 268.80, 283.62, and 199.50 Mbps, respectively. *FTP* has a 1.35x gain in median throughput over the baseline and closely follows the optimal performance.

E. Beam-training Overhead

In this section, we investigate the *beam-training overhead* of *FTP*, *ACO*, and *UbiG* in both single-user (SU) and multi-user (MU) scenarios. The *beam-training overhead* includes both *probing overhead* and *computational overhead*. We calculate the probing latency based on the number of probes required by each approach, following the MAC protocol of IEEE 802.11ad. We conservatively assume zero collisions to demonstrate basic performance, but note that *FTP* is expected to outperform *ACO* even in the presence of collisions due to its lower number of required probes. We measure the runtime of each approach on an Intel i7-7700 CPU to evaluate their computational overhead. Our comparison of performance is not affected by the experimental platform, as the relative gain in computation time is hardware-independent and depends on algorithmic complexity. Finally, we use insights and data from experiments with a 32-element antenna array to conduct simulations with larger antenna arrays, demonstrating the improved gain of *FTP*.

In Fig. 7a, we present the *probing overhead* results for SU scenarios with various antenna array sizes. The *probing overhead* scales logarithmically and linearly for *FTP* and *ACO* while *UbiG* requires a fixed number of probes irrespective

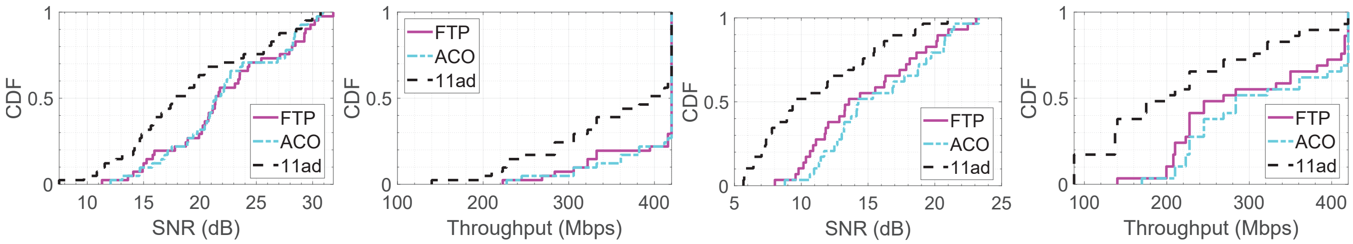

(a) LOS SNR (b) LOS Throughput (c) NLOS SNR (d) NLOS Throughput

Fig. 6: *FTP*'s SNR and throughput performance closely follow the optimal in both LOS and NLOS scenarios.

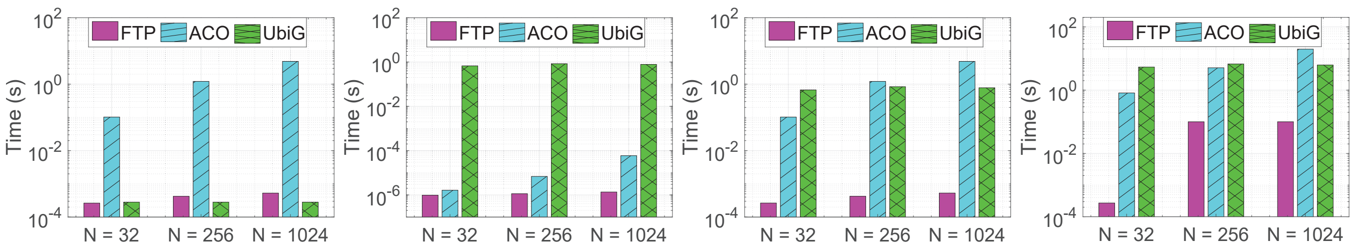

(a) SU probing (b) SU computational (c) SU beam-training (d) MU beam-training

Fig. 7: *FTP* outperforms state-of-the-arts by orders of magnitude in overall *beam-training overhead*. (a) The *probing overhead* scales logarithmically and linearly with the antenna array size N for *FTP* and *ACO* while *UbiG* requires a fixed number of probes. (b) *UbiG*'s algorithmic complexity is considerably higher than *FTP* and *ACO*, resulting in significant *computational overhead*. (c) *FTP* is at least 3 orders of magnitude faster than *ACO* and *UbiG* in overall *beam-training overhead* for SU. (d) With 8 clients, *FTP* is still 2 orders of magnitude faster than *ACO* and *UbiG* in overall *beam-training overhead*.

of the antenna array size N. For *ACO*, such high *probing overhead* prevents it from scaling to large-scale antenna arrays.

Moving on to Fig. 7b, we show the results of the *computational overhead* analysis for SU scenarios with various antenna array sizes. Both *FTP* and *ACO* demonstrate low *computational overhead* as the array size increases due to their low algorithmic complexity. In contrast, *UbiG*'s use of a genetic algorithm results in high algorithmic complexity and significant *computational overhead*, which prevents it to be even deployed on COTS mmWave devices with limited computing resources.

Fig. 7c shows the overall *beam-training overhead* results for SU with different antenna array sizes. As shown in the figure, with increasing antenna array size, *FTP* maintains low *probing overhead* and *computational overhead*, resulting in significant gains in overall *beam-training overhead* compared to *ACO* and *UbiG*. Specifically, for the 32-element antenna array, *FTP* outperforms *ACO* and *UbiG* by up to 3 orders of magnitude. When the antenna array size increases to 1024, the gain of *FTP* increases to 4 orders of magnitude. This demonstrate that *FTP* is a highly-scalable beam-training solution that can find the globally-optimal beam with minimal beam-training overhead.

In Fig. 7d, we present the *beam-training overhead* results for 8 clients with various antenna array sizes. As the number of clients increases, *FTP* requires more than one beacon interval to train all clients with larger antenna arrays, resulting in a significant increase in the *beam-training overhead*. Despite this, *FTP* still outperforms *ACO* and *UbiG* by two orders of magnitude for a 1024-element antenna array. Moreover, for larger antenna arrays, the training time required for one client with *ACO* and *UbiG* is longer than training 8 clients with *FTP*, which shows that *FTP* can handle large-scale antenna arrays and multi-user scenarios simultaneously.

VII. Conclusion

In this paper, we presented *FTP*, a highly-scalable approach for efficient and accurate beam-training in mmWave communications. *FTP* can find the globally-optimal beam with minimal *beam-training overhead* by efficiently estimating per-path direction along with its complex gain. We addressed several challenges, including efficient and accurate path direction estimation, reliable path separation in the CIR, and coherent complex gain for each path. We demonstrated the effectiveness of *FTP* on a mmWave experimental platform with 32 antenna elements and showed that *FTP* achieves optimal SNR performance comparable with state-of-the-art while reducing the *beam-training overhead* by 3 orders of magnitude. Furthermore, we showed that the gain of *FTP* can be even more significant for larger antenna arrays with up to 1024 elements. In summary, *FTP* can efficiently and accurately find the globally-optimal beam, making it feasible to maximize SNR in real-world mmWave communications.

Acknowledgment

This work is supported in part by the NSF under grants CNS-2112471, CNS-2007581 and ECCS-2128567.

References

[1] "Iso/iec/ieee international standard for information technology–telecommunications and information exchange between systems–local and metropolitan area networks–specific requirements-part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications amendment 3: Enhancements for very high throughput in the 60 ghz band (adoption of ieee std 802.11ad-2012)," *ISO/IEC/IEEE 8802-11:2012/Amd.3:2014(E)*, pp. 1–634, 2014.

[2] 3GPP, "5g nr release 16." https://www.3gpp.org/release-16, 2019.

[3] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "Wifi-assisted 60 ghz wireless networks," in *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, 2017, pp. 28–41.

[4] M. Polese, F. Restuccia, and T. Melodia, "Deepbeam: Deep waveform learning for coordination-free beam management in mmwave networks," in *Proceedings of the Twenty-second International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*, 2021, pp. 61–70.

[5] B. Salehi, J. Gu, D. Roy, and K. Chowdhury, "Flash: Federated learning for automated selection of high-band mmwave sectors," in *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 2022, pp. 1719–1728.

[6] H. Yan, B. W. Domae, and D. Cabric, "Mmrapid: Machine learning assisted noncoherent compressive millimeter-wave beam alignment," *arXiv preprint arXiv:2007.12060*, 2020.

[7] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive millimeter-wave sector selection in off-the-shelf ieee 802.11 ad devices," in *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies*, 2017, pp. 414–425.

[8] T. Woodford, X. Zhang, E. Chai, K. Sundaresan, and A. Khojastepour, "Spacebeam: Lidar-driven one-shot mmwave beam management," in *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, 2021, pp. 389–401.

[9] 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*, 2017, pp. 42–55.

[10] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, "Steering with eyes closed: mm-wave beam steering without in-band measurement," in *2015 IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2015, pp. 2416–2424.

[11] Y. Zhang, K. Patel, S. Shakkottai, and R. W. H. Jr, "Side-information-aided noncoherent beam alignment design for millimeter wave systems," in *Proceedings of the Twentieth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, 2019, pp. 341–350.

[12] 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 19)*, 2019, pp. 783–800.

[13] 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 24th Annual International Conference on Mobile Computing and Networking*, 2018, pp. 241–255.

[14] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "Towards scalable and ubiquitous millimeter-wave wireless networks," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, 2018, pp. 257–271.

[15] V. B. Gantovnik, Z. Gurdal, L. T. Watson, and C. M. Anderson-Cook, "Genetic algorithm for mixed integer nonlinear programming problems using separate constraint approximations," *AIAA journal*, vol. 43, no. 8, pp. 1844–1849, 2005.

[16] S. Sur, X. Zhang, P. Ramanathan, and R. Chandra, "{BeamSpy}: Enabling robust 60 {GHz} links under blockage," in *13th USENIX symposium on networked systems design and implementation (NSDI 16)*, 2016, pp. 193–206.

[17] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan, "60 ghz indoor networking through flexible beams: A link-level profiling," in *Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, 2015, pp. 71–84.

[18] H. Hassanieh, O. Abari, M. Rodriguez, M. Abdelghany, D. Katabi, and P. Indyk, "Fast millimeter wave beam alignment," in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, 2018, pp. 432–445.

[19] E. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006.

[20] 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.

[21] D. Ramasamy, S. Venkateswaran, and U. Madhow, "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.

[22] M. E. Rasekh, Z. Marzi, Y. Zhu, U. Madhow, and H. Zheng, "Noncoherent mmwave path tracking," in *Proceedings of the 18th International Workshop on Mobile Computing Systems and Applications*, 2017, pp. 13–18.

[23] A. V. Oppenheim and R. W. Schafer, "Digital signal processing(book)," *Research supported by the Massachusetts Institute of Technology, Bell Telephone Laboratories, and Guggenheim Foundation. Englewood Cliffs, N. J., Prentice-Hall, Inc., 1975. 598 p*, 1975.

[24] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.

[25] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.1," http://cvxr.com/cvx, Mar. 2014.

[26] D. Zhang, M. Garude, and P. H. Pathak, "mmchoir: Exploiting joint transmissions for reliable 60ghz mmwave wlans," in *Proceedings of the Eighteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, 2018, pp. 251–260.

[27] S. Wang, J. Huang, X. Zhang, H. Kim, and S. Dey, "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*, 2020, pp. 1–14.

[28] I. K. Jain, R. Subbaraman, and D. Bharadia, "Two beams are better than one: towards reliable and high throughput mmwave links," in *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, 2021, pp. 488–502.

[29] A. Loch, H. Assasa, J. Palacios, J. Widmer, H. Suys, and B. Debaillie, "Zero overhead device tracking in 60 ghz wireless networks using multi-lobe beam patterns," in *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies*, 2017, pp. 224–237.

[30] R. Zhao, T. Woodford, T. Wei, K. Qian, and X. Zhang, "M-cube: A millimeter-wave massive mimo software radio," in *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, 2020, pp. 1–14.

[31] W.-H. Chen, X. Liu, K. Srinivasan, and S. Parthasarathy, "Fast and optimal beam alignment for off-the-shelf mmwave devices," in *Proceedings of the Int'l ACM Symposium on Mobility Management and Wireless Access*, 2023, pp. 115–123.

[32] M. K. Haider, Y. Ghasempour, D. Koutsonikolas, and E. W. Knightly, "Listeer: Mmwave beam acquisition and steering by tracking indicator leds on wireless aps," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, 2018, pp. 273–288.

[33] M. K. Haider and E. W. Knightly, "Mobility resilience and overhead constrained adaptation in directional 60 ghz wlans: protocol design and system implementation," in *Proceedings of the 17th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, 2016, pp. 61–70.

[34] Y. Chen, Y. Huang, C. Li, Y. Thomas Hou, and W. Lou, "Turbo-hb: A novel design and implementation to achieve ultra-fast hybrid beamforming," in *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*, 2020, pp. 1489–1498.

[35] O. El Ayach, S. Rajagopal, S. Abu-Surra, Z. Pi, and R. W. Heath, "Spatially sparse precoding in millimeter wave mimo systems," *IEEE transactions on wireless communications*, vol. 13, no. 3, pp. 1499–1513, 2014.

[36] Y. Ghasempour, M. K. Haider, C. Cordeiro, D. Koutsonikolas, and E. Knightly, "Multi-stream beam-training for mmwave mimo networks," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, 2018, p. 225–239.

[37] J. M. Steele, *The Cauchy-Schwarz master class: an introduction to the art of mathematical inequalities*. Cambridge University Press, 2004.

[38] A. Networks, "Airfide–a 5g company," https://airfidenet.com, 2018.