

AUTOMS: Automated Service for mmWave Coverage Optimization using Low-cost Metasurfaces

Ruichun Ma*
Yale University
ruichun.ma@yale.edu

Shicheng Zheng*
University of Science and
Technology of China
sc7543@mail.ustc.edu.cn

Hao Pan†
Microsoft Research Asia
panhao@microsoft.com

Lili Qiu
Microsoft Research Asia
UT Austin
liliqiu@microsoft.com

Xingyu Chen‡
University of California San Diego
xic063@ucsd.edu

Liangyu Liu‡
University of Science and
Technology of China
lly@mail.ustc.edu.cn

Yihong Liu‡
University of Glasgow
y.liu.6@research.gla.ac.uk

Wenjun Hu
Yale University
wenjun.hu@yale.edu

Ju Ren
Tsinghua University
renju@tsinghua.edu.cn

ABSTRACT

mmWave networks offer wide bandwidth for high-speed wireless communication but suffer from limited range and susceptibility to blockage. Existing coverage provisioning solutions not only incur high costs but also require significant expert knowledge and manual efforts. In this paper, we present AUTOMS, an automated service framework to optimize mmWave coverage by strategically designing and placing low-cost passive metasurfaces. Our approach consists of three key components: (1) joint optimization of metasurface phase configurations and placement as well as access point beamforming codebooks. (2) a fast 3D ray-tracing simulator for accelerated large-scale metasurface channel modeling. (3) a metasurface design amenable to ultra-low-cost hot stamping fabrication, featuring high reflectivity, near 2π phase control, and wideband support. Simulation and testbed experiments show that AUTOMS can increase the median received signal strength by 11 dB in target rooms and over 20 dB at previous blind spots, and improve the median throughput by over 3× in real-world scenarios.

CCS CONCEPTS

• **Hardware** → **Wireless devices**; • **Networks** → **Wireless access points, base stations and infrastructure**.

*Ruichun Ma and Shicheng Zheng contribute equally to this research

†Hao Pan is the corresponding author

‡Ruichun Ma, Shicheng Zheng, Xingyu Chen, Liangyu Liu, and Yihong Liu did this work as interns at Microsoft Research Asia

KEYWORDS

Millimeter-wave, Coverage optimization, Cloud service, Metasurfaces, Ray-tracing simulator

ACM Reference Format:

Ruichun Ma, Shicheng Zheng, Hao Pan, Lili Qiu, Xingyu Chen, Liangyu Liu, Yihong Liu, Wenjun Hu, and Ju Ren. 2024. AUTOMS: Automated Service for mmWave Coverage Optimization using Low-cost Metasurfaces. In *International Conference On Mobile Computing And Networking (ACM MobiCom '24), September 30–October 4, 2024, Washington, DC, USA*. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3636534.3649347>

1 INTRODUCTION

The millimeter wave (mmWave) frequency band is critical for high-speed wireless communication in 5G and beyond due to its large bandwidth [1, 19]. However, mmWave signals suffer from limited range and blockage due to their high frequency [52], which hinders wide adoption. One potential solution is to utilize massive phased antenna arrays (*e.g.*, \$3000+ [46]) and create narrow beams to increase the received signal power and extend the line-of-sight (LOS) range [6]. To cover multiple non-line-of-sight (NLOS) areas in a complex deployment environment, a common solution is to deploy multiple mmWave access points (APs) or base station [29]. These solutions are not only costly (*e.g.*, around \$300 for each AP [4]), but also lack scalability due to the considerable domain knowledge and manual labor required for deployment. This is because the APs must be strategically positioned and configured to prevent interference and guarantee coverage.

Smart surfaces or metasurfaces, composed of specially designed sub-wavelength elements, have powerful capability to manipulate electromagnetic (EM) wave propagation and

enhance mmWave coverage beyond what simple mirror reflectors [8, 62] offer. Unbound by the law of reflection, they enable flexible wavefront control, such as reflecting at various angles in 3D space, and maintain high performance with minimal surface area, as demonstrated in Figure 13. Previous works [5, 9, 13, 15, 26, 31, 59] design programmable surfaces to enhance coverage in blind spots and notably improve channel conditions. However, these solutions often rely on tunable circuit components and computing devices to enable rapid beam switching, resulting in costly hardware, considerable computational overhead, and increased power consumption. MilliMirror [41] uses passive metasurfaces to lower cost as a proof of concept, but handles only limited scenarios – beam reflection towards a specific direction by manually identifying the placement of metasurfaces. It remains open how to unleash the full power of metasurfaces, balancing trade-offs between effectiveness, complexity, cost, and deployment burden on users.

In order to realize the potential of metasurfaces, there are several significant challenges. First, existing solutions require substantial expertise in wireless signal propagation for guiding their design and deployment. Manual configuration is error-prone and often results in less-than-ideal performance [32]. Second, automating optimization hinges on a reliable wireless channel model that accurately captures EM wave propagation in complex scenarios. Existing simulators, like Wireless Insite [44], fall short by being either slow or offering limited metasurface support, such as only enabling linear phase configurations [43]. Developing a fast and accurate simulator that supports general metasurface operations is a challenging problem. Third, metasurface-based methods must be considerably more cost-effective than alternatives to be appealing. Current strategies do not demonstrate a substantial economic advantage over multi-AP solutions (Table 1). This necessitates the creation of an affordable and efficient passive metasurface design.

We propose AUTOMS, an automated service framework that optimizes mmWave coverage for given areas, such as apartments or office floors, using strategically designed and placed low-cost passive reflective metasurfaces, as depicted in Figure 1. These metasurfaces alter mmWave propagation to overcome blockages and concentrate power toward desired regions. To minimize costs and simplify deployment, we employ passive metasurfaces that are fixed post-fabrication. Moreover, we leverage the AP’s dynamic beam steering capability together with passive metasurfaces to serve different areas and handle user mobility.

Our system consists of following three major components:

Automated optimization: To automate metasurface-assisted mmWave coverage planning, we develop a novel joint optimization framework. It determines optimal network deployment configurations, including the phase settings and

Method	Type	Coverage	Cost	Hardware count
AP [4, 54]	Active	Manual	≈ \$300	per AP
Relay [2]	Active	Manual	≥ \$200	per relay
mmWall [11]	Varactors	Manual	≈ \$10000	28 × 76
milliMirror [41]	Passive	Manual	≈ \$15	80 × 80
AUTOMS	Passive	Auto	≈ \$1	160 × 200

Table 1: Comparison between mmWave coverage solutions. We show cost details in section 4

placement of one or more metasurfaces, AP(s)’ codebook (phased array configurations), and, if necessary, the positioning of the APs. We iteratively identify and optimize hyper-configurations, such as the quantity, sizes, and placement of metasurfaces. Each iteration refines the phase settings of metasurfaces and APs through gradient descent. While our framework primarily aims to maximize the aggregate channel capacity, it is also flexible enough to support alternative objectives and may extend to active or programmable metasurfaces.

Fast simulator: We develop a fast 3D ray tracing simulator capable of computing channel response matrices for metasurfaces with tens of thousands of elements, utilizing both software and hardware acceleration techniques. Our simulator processes approximately 1.6 billion wireless propagation paths in 3 minutes using an NVIDIA A100 GPU, supporting large metasurfaces by element count. To our knowledge, no current ray-tracing simulators can handle large metasurfaces with arbitrary phase configurations due to the prohibitive computational costs.

Low-cost metasurface: We design a passive metasurface to achieve several important properties: high reflectivity ($\geq 90\%$) to minimize energy loss, near 2π phase range for wavefront control, wideband support (over 8 GHz) on mmWave bands, and compatibility with low-precision but cost-effective hot stamping. By employing surface substrates, *i.e.*, paper and PVC plastic, we induce specific phase delays, forming an impedance transformer and thus “amplifying” phase shifts of different metallic patterns. With this approach, we achieve the desired properties using single-layer split-ring patterns that can be hot stamped onto paper, allowing for virtually cost-free metasurface production.

We implement AUTOMS as a cloud-based service featuring a user-friendly, multi-step workflow illustrated in Figure 2. Users can employ commercial applications, such as Polycam [39], to capture their deployment environment using a mobile device, resulting in a 3D layout model. Concurrently, we design a straightforward yet effective technique to infer reflection ratios of walls and objects in the environment. This is achieved by collecting Received Signal Strength (RSS) measurements during the scanning process and correlating them with RSS predictions from ray tracing, enabling automatic estimation of both the 3D model and the wall reflectivity. Subsequently, AUTOMS leverages this data to compute optimized deployment configurations, metasurface designs, and

an updated codebook for the AP to work together with the metasurfaces.

When a cable/DSL company installs Internet service, they can offer a value-added service by capturing the 3D layout and uploading it to the AUTOMS platform to generate metasurfaces and provide placement instructions. Customers facing mmWave Wi-Fi coverage issues can also submit their 3D layout to the AUTOMS online service for assistance. We introduce a novel concept in network deployment – *surfaces as a service*. This approach transfers the complexity of hardware design to the cloud service, aligning with trends seen in 5G network development.

We evaluate AUTOMS using a public 3D layout dataset [18] and through real-world testing. Simulations show that our system markedly enhances mmWave coverage across various scenarios, achieving a median RSS gain of 12.1 dB compared to a single-AP setup. Field tests further confirm that our method boosts median RSS by 11 dB in target areas and over 20 dB in blind spots, elevating median throughput from 77 Mbps to 373 Mbps. Our metasurfaces surpass mirror reflectors four times their size by a median of 7 dB in RSS improvement. As discussed in subsection 5.4, AUTOMS proves resilient to environmental changes. By restricting the vertical search space for placements, we prioritize above-head locations for APs and metasurfaces, reducing the likelihood of human movement disrupting the links. Combined with dynamic AP beam steering, AUTOMS ensures dependable coverage regardless of environmental dynamics.

To summarize, we make the following contributions:

- We present AUTOMS, the first automated service framework for metasurface-aided mmWave network deployment.
- We build an efficient 3D ray-tracing wireless channel simulator with a series of acceleration techniques. We open-source the code to the research community at <https://github.com/microsoft/AutoMS>.
- We design novel hot-stamped passive metasurfaces for coverage enhancement. It achieves high reflectivity, near 2π phase control, and wideband operation, with orders of magnitude lower cost than prior solutions (Table 1).
- We evaluate AUTOMS extensively in 5 simulated environments and 3 real-world environments, achieving robust mmWave coverage in diverse deployment scenarios.

2 RELATED WORK

2.1 mmWave Coverage Extension

Conventional Methods Passive mirror reflectors are used to extend mmWave coverage in data centers [62], indoor [23], and outdoor settings [37], by reflecting signals around obstacles into NLOS areas. However, they are limited by the law of reflection, where the reflection angle is tied to the incident angle and the orientation of the reflector. As illustrated in

Figure 13, even optimized placement of traditional reflectors often fails to achieve complete mmWave coverage in complex environments. Some studies [22, 56] attempt to improve link reliability by utilizing pre-existing large reflectors, but this is generally only practical in certain scenarios. Others [2, 53] introduce active mmWave relays that re-transmit signals to widen coverage and bolster link resilience. Despite their effectiveness, such active solutions are costly and power-intensive, akin to adding extra APs.

Reconfigurable/Active surfaces. Several existing works focus on developing reconfigurable surfaces with active elements that can dynamically modulate wireless signals. A recent series of research presents end-to-end systems aimed at manipulating signal paths to enhance receiver channel conditions [5, 9, 13, 15, 26, 31, 57–59]. However, these systems, tailored for sub-6 GHz Wi-Fi frequencies, are not directly transferable to mmWave communications due to differences in channel modeling, endpoint architecture, and application scenarios. Compared to lower frequencies, mmWave has a limited range and is more susceptible to blockage. Metasurfaces have emerged as a promising solution [10, 11, 20, 25]. However, programmable metasurfaces face several challenges. They require costly tunable circuits and printed circuit board (PCB) substrates. While individual elements are inexpensive, a metasurface with thousands of them can become more expensive than mmWave APs, diminishing the appeal of using such surfaces. Additionally, they rely on external controllers for wiring and power, adding complexity to the control system and installation. The interaction and feedback loop between an active surface and the endpoints remain unresolved issues.

Passive metasurfaces. MilliMirror [41] introduces a cost-effective, 3D-printed reflective metasurface with beam synthesis capabilities, although its production process is not yet mainstream. While showcasing several deployment case studies, MilliMirror lacks a systematic method for automating metasurface design and placement and requires significant expertise. It, along with other reconfigurable intelligent surface (RIS) research, aims to enhance single links, whereas our work seeks to optimize network-wide performance. PM-Sat [35] offers a unique passive metasurface design targeting the optimization of individual terrestrial and satellite links. Additionally, recent studies [47, 48] have utilized hot-stamped passive metasurfaces for security applications, but these designs suffer from a power loss of over 7 dB. In contrast, AUTOMS pioneers the use of hot stamping in metasurfaces to boost SNR with minimal power loss.

Our approach. We develop a comprehensive approach that inputs a 3D scene model and determines the best placement and settings for metasurfaces and AP codebooks. We employ high-performance passive metasurface designs to ease production, deployment, and to lower costs. Despite its passive




Figure 1: Metasurfaces create paths to circumvent blockages (beam 1) and extend coverage (beams 2&3) in the 3D environment

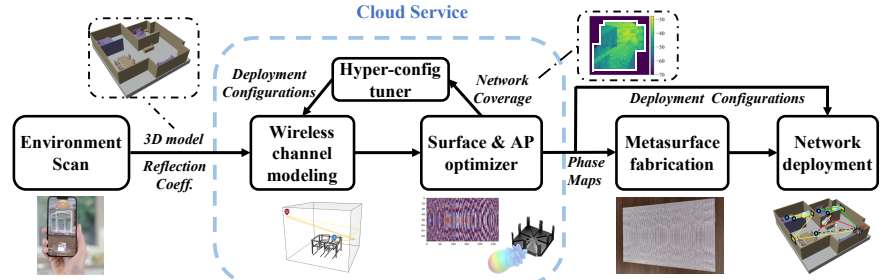


Figure 2: AUTOMS overview. Based on an environmental scan, our service framework outputs the optimized deployment configurations for passive metasurfaces and AP(s) tailored to the target 3D model. After surface fabrication, we deploy the metasurfaces and AP accordingly

nature, our metasurface can accommodate endpoint mobility and environmental shifts by collaborating with the AP’s phased array, which selects suitable beamforming codewords.

2.2 Ray Tracing-based Channel Modeling

Accurate simulation of wireless signal propagation and channel modeling is crucial for designing and deploying wireless networks. Ray tracing (RT) methods, which approximate the Maxwell equations at high frequencies (*e.g.*, as optical rays), represent the field as a set of rays that reflect, diffract, and scatter through the environment. While several RT simulators like WinProp [3], Aster [16], CrossWave [17], and Wireless Insite [44] are available for commercial or academic use, only the latest Wireless Insite [43] version includes metasurface support. However, based on our discussions with their technical team, it only allows for reflection towards a fixed angular range and lacks the capability to model individual metasurface elements, precluding optimization of the elements themselves.

Prior PIS research [27, 28] relied on a basic ray-tracing simulator that only offered control at the surface level and assumed environments without obstacles. Recently, some researchers proposed open-source RT simulators, like Opal [14], SimRIS [7], and NeRF2 [61]. Opal, fast due to its OptiX API implementation, unfortunately lacks metasurface modeling capabilities [14]. SimRIS, a MATLAB-based simulator, can model channels for RIS-assisted MIMO systems but is not equipped for handling large metasurfaces [7]. NeRF2 uses a neural network for channel modeling but does not include metasurface support and necessitates considerable training for each scenario. Considering that a metasurface can have tens of thousands of elements, creating a multitude of signal paths, scalability becomes a significant challenge. To address this, we have developed a 3D mmWave ray tracing simulator capable of rapidly and accurately modeling wireless channels for general metasurfaces.

3 AUTOMS DESIGN

3.1 Overview

Considering the symmetry between upstream and downstream links in mmWave communications, our study focuses on a downstream scenario specifically, transmission from the AP to the receiver in an indoor setting. We presume the mmWave AP has a phased array capable of beam steering and that reflective metasurfaces are installed on the walls. In line with the IEEE 802.11ad standard, we utilize a beam scanning process where the AP sequentially transmits using all codewords. The receiver then selects and reports back the codeword that yields the highest SNR. Therefore, the achievable link for the receiver is the best link among all codewords.

Figure 2 illustrates the workflow of our system. Using RoomPlan [12] on a mobile device (section 4), we first create a 3D model of the environment for network deployment, serving as the input to AUTOMS. Subsequently, our framework iteratively executes the following steps: (1) Our proposed ray-tracing simulator generates wireless channel response matrices between the AP (TX), metasurfaces, and potential receiver (RX) locations based on the 3D model, influenced by the deployment’s hyper-configurations such as metasurface and AP placements. (2) A phase map optimizer employs gradient descent to fine-tune the metasurfaces’ phase settings and APs’ codebooks to maximize a network-wide metric, like sum capacity, and provides feedback on the optimized performance to the hyper-configuration tuner. (3) The hyper-configuration tuner, typically using simulated annealing, updates the deployment hyper-configurations, prompting a return to Step 1 for updates. (4) The process concludes when the optimized performance stabilizes, yielding the final phase maps and deployment configurations for the actual metasurface fabrication and installation.

3.2 Optimization Framework

Our optimization framework comprises a hyper-configuration tuner and a phase map optimizer for mmWave network deployment. Hyper-configurations are high-level deployment

parameters, such as the quantity, dimensions (element count), and locations of metasurfaces, as well as the orientation(s) and location(s) of AP(s). The phase map optimizer operates as the inner loop, determining the phase configurations for the given hyper-configurations.

Hyper-configuration tuner. AUTOMS allows for the optimization of certain deployment configurations based on the specific scenario and user preferences, while users can manually set the remaining according to their requirements. For instance, if the AP's position is predetermined, AUTOMS focuses on optimizing the other parameters. Various algorithms can optimize these configurations; our implementation employs simulated annealing, which has proven effective for our purposes. The hyper-configuration tuner, informed by the phase map optimizer's performance metrics, iteratively explores and updates configurations until performance stabilizes (*i.e.*, no improvements are observed in the last iterations). The resulting final configurations guide the fabrication and deployment of metasurfaces.

Phase map optimizer. The phase map optimizer identifies the optimal phase configurations for metasurfaces and mmWave AP phased arrays to maximize the coverage objective function in a target environment. This optimization is based on the channel matrices provided by the simulator and the current hyper-configurations. It determines the phase shift values for each metasurface element, denoted as $W_{ms}(k)$ with k representing the metasurface index, and selects a set of beamforming codewords for the APs, indicated by $W_{ap}(j)$ where j is the codeword index. Upon achieving convergence, these phase configurations are discretized for practical deployment.

Objective function. There are many possible metrics that can quantify the coverage of an area. Our design is intentionally decoupled from physical layer specifics, such as modulation schemes or hardware implementations, enabling our system's compatibility with diverse mmWave devices and protocols. Consequently, we select the sum of channel capacities between the transmitter (TX) and all potential receiver (RX) locations as our optimization goal, though our framework is versatile enough to accommodate alternative objectives. As outlined in subsection 3.3, channel matrices involving TX, RX, and metasurfaces are computed using ray tracing. Our optimization problem is differentiable, allowing us to employ gradient descent with the Adam optimizer to enhance sum capacity. Conceptually, this optimization can be viewed as the strategic distribution of power from the metasurfaces to all RX locations.

We define the j th codeword among all N_c code words as $W_{ap}^{(j)} \in \mathbb{C}^{1 \times N_t}$ and the phase configuration of k th surface among all M surfaces as $W_{ms}^{(k)} \in \mathbb{C}^{N_{ms} \times 1}$, where N_{ms} and N_t are the numbers of metasurface elements and transmitting antennas, respectively. Each value in the phase configuration and

codebook matrix is a phase shift without amplitude change, denoted as $e^{j\theta}$, where θ is the phase we aim to determine. The sum of channel capacity between the TX and RX locations is used as the optimization objective and derived as follows:

$$\max_{W_{ms}, W_{ap}} \sum_{i=1}^{N_r} \log_2 \left(1 + \frac{S_i(W_{ms}, W_{ap})}{Noise} \right) \quad (1)$$

where N_r denotes the number of RX locations and S_i denotes the received signal at the i th receiving antenna as a function of metasurface phase maps and AP codebook. We set a noise floor of -70 dBm to match our experiment hardware (*i.e.*, 0 capacity if RSS is below -70 dBm). Next, we describe the calculation for received signal strength S_i . Let P_t denote the transmission power, G_t denote the transmission antenna gain, G_r denote the receiver antenna gain. Since the AP can switch between multiple codewords, we take the maximal achievable signal strength from all codewords as S_i :

$$S_i(W_{ms}, W_{ap}) = \sqrt{P_t G_t G_r} \max_{j \in [1, N_c]} W_{ap}^{(j)} (\mathbf{H}_{ms,i} + \mathbf{H}_{T-R_i}) \quad (2)$$

where $\mathbf{H}_{ms,i} \in \mathbb{C}^{N_r \times 1}$ is the channel matrix characterizing the propagation paths affected by the metasurfaces between the TX and the i th RX, $\mathbf{H}_{T-R_i} \in \mathbb{C}^{N_r \times 1}$ is the channel matrix between the TX and the i th RX, characterizing the propagation paths not traversing any metasurfaces. Due to a small N_c , S_i 's optimal value can be obtained easily without an optimization solver. The metasurface-affected channel matrix, $\mathbf{H}_{ms,i}$ is a function of $\{W_{ms}^{(k)}\}_{k=1}^M$. If we only consider the first-order reflections from metasurfaces, then we have

$$\mathbf{H}_{ms,i} = \sum_{k=1}^M (\mathbf{H}_{T-ms} W_{ms}^{(k)} \mathbf{H}_{ms-R_i}) \quad (3)$$

where $\mathbf{H}_{T-ms} \in \mathbb{C}^{N_t \times N_{ms}}$ is the channel matrix between the transmitting antennas and the metasurface elements, $\mathbf{H}_{ms-R_i} \in \mathbb{C}^{N_{ms} \times 1}$ is the channel matrix between the metasurface elements and the receiver. We further extend the equation above to support reflections between multiple metasurfaces, but we omit the details here.

Thus, we need to know the following channel matrices to compute the objective function: \mathbf{H}_{T-R} , \mathbf{H}_{T-ms} , \mathbf{H}_{ms-ms_j} , and \mathbf{H}_{ms-R} . These channel matrices are calculated using the ray tracing simulator, as described in subsection 3.3.

3.3 Efficient Channel Simulator

To address the optimization problem, we simulate the channel response matrices between the mmWave APs, metasurfaces, and receivers in the environment. Each metasurface element is modeled as a reflective antenna with a phase shift, allowing us to treat different channels similarly to those between source and destination antenna(s). This modeling is independent of




Figure 3: Illustration of path loss when reflected by a reflector vs. a metasurface unit

the metasurface hardware design and assumes no coupling between surface elements.

Ray tracing-based modeling. We employ ray tracing to accurately and efficiently simulate wireless signal propagation, tracing rays from the source and modeling their paths and directions to compute the channel response. Due to the high attenuation of mmWave, we primarily focus on modeling reflections and direct LOS propagation paths.

We calculate the channel response in two stages. Initially, we launch rays to identify valid propagation paths, emitting numerous rays from the sources. Some rays reach the destination antennas, while others are discarded due to missing the destination or low signal strength. We record the propagation paths together with any intersecting reflectors for rays that reach the destination. Next, we compute the channel response for each valid path and aggregate these responses to determine the overall channel response at the destination, encompassing both amplitude and phase information. To validate the accuracy of our channel modeling, we compare our simulation results with real-world measurements. As shown in Figure 4, our simulations correspond closely with real-world measurements. Further validation is conducted in section 5.

Channel response calculation. Given the propagation paths between a source and a destination, we need to calculate the channel response. Here, we describe the channel model used in our simulation. We first analyze the LOS path. The channel response can be derived as $h_{LOS} = \frac{\lambda}{4\pi r} e^{\frac{2\pi r}{\lambda}}$, where λ is the wavelength corresponding to the carrier frequency, and r is the distance from the AP to the user. For NLOS paths that include multi-order reflections from objects, we analyze each reflection separately and multiply them together. When a path goes from the previous point to the next point, two types of reflection can happen as shown in Figure 3. When the reflector is sufficiently large to be considered as a mirror, *e.g.*, walls and floors, such that the previous point is within the near-field distance of it, we have $\theta_{inc} = \theta_{out}$ and the channel response of this reflection is

$$h_{re} = \tau \frac{\lambda}{4\pi(r_{inc} + r_{out})} e^{\frac{2\pi}{\lambda}(r_{inc} + r_{out})} \quad (4)$$




Figure 4: Comparison between simulation results of our ray-tracing based channel modeling and real-world RSS measurements

When the reflector is small enough to be considered as a far-field scatterer, *e.g.*, a metasurface element, we have scattering:

$$h_{re} = \tau \frac{\lambda}{4\pi r_{inc}} e^{\frac{2\pi}{\lambda} r_{inc}} \frac{\lambda}{4\pi r_{out}} e^{\frac{2\pi}{\lambda} r_{out}} \quad (5)$$

where τ denotes the reflection coefficient of the reflector or surface element. We use the equations above iteratively for multi-order reflections. Next, we sum the LOS and NLOS channel responses to calculate the overall channel response for one source-destination pair as follow:

$$h_{s-t} = \sum_{i=1}^P h_{re}^{(i)} + h_{LOS} \quad (6)$$

where P is the number of NLOS propagation paths between the source and destination. By replacing the source and destination with TX, RX, or metasurfaces, we have all the channel response matrices needed for the objective function, *i.e.*, \mathbf{H}_{T-R} , \mathbf{H}_{T-ms} and \mathbf{H}_{ms-R} . We plug into these values into Equation 3 and solve the optimization problem.

Acceleration. We develop several techniques to boost simulation efficiency. Initially, we optimize the propagation path discovery, the most resource-intensive step in ray tracing, by sharing path information among nearby destinations. While casting a large number of rays reduces the likelihood of overlooked paths and heightens accuracy, it also incurs substantial memory and computation costs. To mitigate these costs, we divide the area of interest into a grid system where each grid can exchange propagation path details with its neighbors, allowing for the discovery of paths using significantly fewer rays. Moreover, we refine data storage by cataloging path information through the set of reflectors traversed from source to destination, calculating the cascaded coordinate transformation matrices for mirroring. This approach allows multiple TX-RX pairs in close proximity to reuse these matrices, conserving memory and computation time when calculating channel responses. Finally, we implement our ray tracing using OptiX [36], which capitalizes on RT cores for GPU acceleration. With this, we can simulate approximately 1.6 billion wireless channel responses in just 3 minutes on an NVIDIA A100 GPU.

Material reflection coefficients. The reflection coefficient of materials significantly influences the wireless channel. In our experiments, we use values from existing literature [24, 45] and enhance our database by measuring mmWave properties of uncharted materials. To improve simulation precision, we develop a method that utilizes RSS measurements to deduce material reflection coefficients. During environmental scanning for 3D modeling, a mobile device with mmWave connectivity also gathers RSS data throughout the area. At each location where RSS is collected, our ray tracing simulator identifies the propagation paths of all rays reaching that point. We then treat the reflection coefficients as variables to be optimized, calculating RSS values based on these propagation paths. Starting with coefficients from the literature, we employ gradient descent to find values that minimize the discrepancy between the calculated and actual RSS measurements. As shown in Figure 8, using optimized material properties results in simulations that are more accurate than those based solely on literature values.

3.4 Low-cost Metasurface Design

In this section, we present the development of a reflective and passive mmWave metasurface design that is employed in our proposed AUTOMS.

Design goals. We set the following design goals: (i) high reflectivity (*e.g.*, $\geq 90\%$), to minimize energy loss due to metasurface.; (ii) full 2π phase control, which is crucial for beamforming; (iii) wideband operation, covering a 60 GHz channel bandwidth; and (iv) ultra-low cost, *i.e.*, considerably cheaper than other solutions (see Table 1), achieved through hot-stamping fabrication on common materials.

Challenges and our approach. Achieving 2π phase control and high reflectivity for mmWave with a low-complexity and low-cost design is a challenging task. Previous works have used either meta-atoms with rib-like arrangements [11] or 3D printed materials with varying thicknesses [41], which poses challenges in terms of surface thickness, assembly, and fabrication. In contrast, we adopt 2D metallic patterns on a thin (~ 0.5 mm) dielectric substrate and a metallic sheet to induce reflection phase shifts (Figure 5d), achieving simplicity and sub-mm thickness. Figure 5a shows a split ring-based pattern, which is the template of different patterns for phase control. By varying the pattern dimensions, we can control the phase shift. However, a single layer has a limited control on the phase shift [38]. One solution is to cascade more layers of patterns [35], but this significantly increases design complexity and power loss due to dissipation in the substrate.

Substrate as an impedance transformer. We propose a novel design that leverages the surface substrate between metallic patterns and a metal sheet as a microwave component. As incident waves propagate through the substrate, it

functions like a transmission line, inducing a phase delay. Consequently, we consider substrate thickness to be a critical design parameter. With a phase delay of $\frac{\pi}{2}$, the substrate acts as a quarter-wave impedance transformer, as detailed in *Microwave Engineering* [40], which effectively transforms the zero impedance of a metal sheet into infinite impedance. We control the impedance mismatch between air and this infinite impedance using a metallic pattern, drawing inspiration from prior work [30]. Therefore, a set of split-ring patterns, spanning a limited range of surface pattern admittance, results in a broad reflection phase range.

Equivalent circuit analysis. To enhance the understanding of phase control, we refer to the equivalent circuit model depicted in Figure 5b. Here, air and the surface substrate are represented as transmission lines with impedances denoted by $Z_0 = 376.7 \Omega$ and Z_1 , respectively. The substrate-based transmission line possesses a thickness l and a phase constant β . EM waves arriving at the surface induce current along the pattern, enabling us to model the metallic pattern as a circuit component with an admittance Y_p . This admittance is purely imaginary, as the resistance of the metallic patterns is negligible. The metal sheet provides a termination for the substrate transmission line with an impedance $Z_m = 0$ (short-circuited). Consequently, the substrate and the metal sheet together form a short-circuited transmission line, as described in *Microwave Engineering* [40], resulting in an impedance:

$$Z_L = Z_1 \frac{Z_m + jZ_1 \tan(\beta l)}{Z_1 + jZ_m \tan(\beta l)} = jZ_1 \tan(\beta l) \quad (7)$$

which describes the total impedance of the substrate and the metal sheet (the whole surface except metallic patterns). The phase delay caused by the substrate is around $\frac{\pi}{2}$ as a quarter wave impedance transformer, so $Z_L \approx jZ_1 \tan(\frac{\pi}{2}) = \infty j$, which indicates the substrate transforms the impedance of metal sheet from 0 to ∞j .

Subsequently, the surface can be modeled as two parallel circuit components: (i) a metallic pattern with admittance Y_p , and (ii) a combination of the substrate and metal sheet, which we represent with an impedance Z_L . The equivalent impedance of the entire surface is determined as follows:

$$Z_S = \frac{1}{Y_p + 1/Z_L} \approx \frac{1}{Y_p} \quad (8)$$

We observe that the Z_S undergoes significant variations when the admittance Y_p fluctuates near zero. Following this, the reflection from the surface can be modeled as the interaction between the air transmission line with impedance Z_0 and a load represented by Z_S . Consequently, we can express the reflection coefficient (S_{11}), *i.e.*, the ratio of the incoming signal E_{in} to the reflected signal E_{out} [40], as follows:

$$\Gamma = \frac{E_{in}}{E_{out}} = \frac{Z_S - Z_0}{Z_S + Z_0} = \frac{1/Y_p - Z_0}{1/Y_p + Z_0} \quad (9)$$




Figure 5: AUTOMS metasurface element design. The reflection phase of the surface varies with different pattern admittance in (c); It provides high reflection power and near 2π phase coverage over wideband (e.g., 57-65 GHz)

Note that Y_p is an imaginary number and Z_0 is a real number. Thus, we have $\|\Gamma\| = \frac{\|1/Y_p - Z_0\|}{\|1/Y_p + Z_0\|} = 1$, which means almost all incoming signal power is reflected. The phase of Γ , *i.e.*, phase shift between incoming and reflected signals, becomes:

$$\angle\Gamma = 2 \arctan \frac{1}{\text{Im}(Y_p)Z_0} + \pi \quad (10)$$

By modifying the dimensions of the split-ring pattern, we effectively adjust the pattern's admittance Y_p , for example, from $-0.005j$ S to $0.035j$ S in our design. Upon substituting these values, we achieve a nearly 2π phase shift range, provided that Y_p offers an adequate tuning range.

Selecting patterns. We select a set of split-ring patterns with varying admittances Y_p to span the entire 2π phase range. Through HFSS simulations, we adjusted three ring pattern parameters—radius (R), width (W), and gap (G), as shown in Figure 5a—to obtain a broad admittance spectrum. Figure 5c displays the admittance and corresponding phase shifts for six exemplary patterns. By varying the ring radius from 0.6 mm to 1.1 mm, widths between 0.3 mm and 0.5 mm, and gaps from 0.3 mm to 0.7 mm, we achieved admittance values ranging from $-0.005j$ S to $0.035j$ S. Given our hot-stamping fabrication precision of approximately 0.1 mm, we adjusted the dimensions in increments of 0.1 mm, resulting in a maximum phase error of 30 deg. This error is within acceptable limits for wavefront manipulation and is more precise than many commercial devices; for example, phased arrays in commercial routers typically have only $\frac{\pi}{4}$ granularity [60].

Wideband performance. Selecting patterns with wideband performance is crucial. The metasurface must ensure a high reflection coefficient and consistent phase control across mmWave channels, such as 58-70 GHz for 802.11ad networks. Since a uniform phase offset across all patterns does not alter the metasurface's beam pattern, our focus is on maintaining a steady phase differential between different patterns. The dimensions of the patterns, particularly the ring width, affect both admittance and wideband performance. We chose multiple patterns that maintain high reflection and a consistent phase across 57 GHz to 65 GHz, adequately covering IEEE 802.11ad/ay channels. For precise phase control, we employ

18 patterns in our design. Figure 5d illustrates the performance of six such patterns. We anticipate that our design could be adapted to other frequency bands as well.

Metasurface pattern array. To create the metasurface design, we integrate a set of patterns that facilitate phase control. The patterns are spaced 2.5 mm apart, approximately half of the wavelength, to minimize undesired coupling between different patterns. The optimization framework yields the optimized phase configuration for each element, which we then match with the corresponding patterns to construct the metasurface. The efficacy of our design is demonstrated through beam steering tests, as shown in Figure 10.

Enabling low-cost fabrication. Our ultra-thin single-layer design enables the hot stamping of metallic patterns onto affordable materials without compromising performance. Although standard PCB production offers high precision (up to 0.03 mm), it becomes costly for large metasurfaces—for instance, a $0.5 \times 0.5 \text{ m}^2$ RF PCB may exceed \$1000. Metasurfaces, however, operate differently as RF signals propagate orthogonally through the thin substrate, not along PCB traces. At just 0.5 mm thick, excluding the metal layer, our design permits the use of less expensive substrates such as paper and PVC, which, despite higher loss tangents, dramatically reduce costs. To confirm feasibility, we conducted simulations considering material loss tangents, precise substrate thickness, and small fabrication-induced discrepancies. The results, depicted in Figure 5d, reveal that our design's power loss is minimal, only 0.3 dB higher than that of a Rogers PCB substrate. The detailed fabrication process is outlined in section 4. While recent studies [21, 48] have also employed hot stamping, their designs solely affect cross-polarization phase and suffer from over 7 dB power loss, rendering them unsuitable for our goal of coverage enhancement.

3.5 Practical Considerations

Robustness to changes. Our flexible optimization framework employs several strategies to handle changes: (i) The AP can pick the optimal beam steering codeword to dynamically counter minor location variations, *e.g.*, within < 20 cm.




Figure 6: Fabrication process of our proposed ultra-low-cost passive metasurfaces

(ii) For substantial changes in AP location or furniture rearrangement, we can adjust the AP codebook or metasurface placement while keeping the metasurface configuration static. (iii) To minimize disruptions from human activity, we position the AP and metasurface above typical human height during optimization (see Figure 18 and Figure 19). (iv) We account for various situations, optimizing for either average or worst-case performance, *e.g.*, different door states (refer to Figure 20).

3D modeling accuracy. We employ 3D Lidar mapping and image classification to scan environments, identifying objects like wood tables, displays, walls, and doors, and determining their material types to assign appropriate reflection coefficients (detailed in section 4). To improve accuracy, we refine these coefficients with RSS measurements (see subsection 3.3 and Figure 8). Our optimization is resilient to minor inaccuracies, such as wall shifts at the centimeter level, maintaining coverage integrity. Continuous RSS monitoring can also detect and adjust to environmental changes over time.

Optimization customization and scalability. Our framework supports multiple optimization objectives, defaulting to sum capacity. Utility weights can be assigned to different regions, allowing customization of the objective function for varied use cases. It also supports optimization for single or multiple AP placements with metasurfaces by including additional variables. We concentrate on single AP configurations with metasurfaces to effectively address coverage in blind spots. The scalability of our method depends on the GPU memory required for channel matrix calculations in the ray-tracing simulator; enhancing this calculation’s efficiency will boost our approach’s scalability.

4 IMPLEMENTATION

Modeling 3D environment. Mobile devices, *e.g.*, iPhones and iPads, are equipped with a Lidar sensor to provide 3D depth information. RoomPlan [12] is an API by Apple that runs a machine learning algorithm on both Lidar and RGB camera information to generate 3D room plan information.

Based on this feature, many mobile apps have been developed to create 3D models of real-world environments, *e.g.*, polycam [39]. In this paper, we use off-the-shelf solutions to implement 3D model and mainly focus on fast simulation of EM wave propagation in a given environment and effective optimization of metasurface design and placement.

Software implementation. We implement AUTOMS as a software running on an Azure cloud server with one NVIDIA A100 GPU. The computation hardware cost is amortized and very low as a cloud instance. We leverage an existing toolkit [33] for the algorithm used in hyper-config tuner. We use Optix library [36] to implement wireless channel simulator and performs phase map optimization using gradient descent in pyTorch.

Metasurface fabrication. Our metasurfaces are constructed using common low-cost materials, detailed in Figure 6. Initially, a laser printer creates the pattern. Then, an aluminum hot-stamping foil is placed over the print and fed through a laminator at approximately 170°C, causing the metallic powder to bond to the ink. Once the foil is removed, the metallic patterns remain on the paper. To complete the assembly, we layer the materials: office printer paper (0.1 mm) holds the metallic pattern, and a plastic sheet (0.3 mm) augments the substrate thickness. These are combined with a metal ground plane using double-sided tape, as illustrated in Figure 5a.

Cost comparison. Existing mmWave coverage enhancement solutions are priced similarly to or above that of a mmWave AP, diminishing their appeal (Table 1). For instance, mmWall features 76×28 elements with two MAVR-000120-14110P varactors each, and at \$2.4 per varactor, the material cost for varactors alone reaches \$10000. Substituting varactors with less expensive PIN diodes might reduce costs below \$1000, yet considerable engineering effort is required for comparable performance. Passive metasurfaces present a more affordable alternative. A $10 \times 10 \text{ cm}^2$ MilliMirror sample with 80×80 elements is priced at about \$15, making a $40 \times 50 \text{ cm}^2$ surface approximately \$300, requiring specialized 3D printing resources. In contrast, our AUTOMS metasurface of the same size with 160×200 elements costs merely \$3, utilizing readily available materials and tools. This cost accounts for hot stamping foils (\$0.1), printer paper (\$0.1), a PVC sheet (\$1), and a 1 mm thick metal plate (\$2). Mass production could further reduce costs. Additionally, while we use a thick metal plate for stability and ease of manual fabrication, a thin 0.05 mm metal sheet suffices for functionality at a significantly lower cost, bringing the total metasurface cost to around \$1. Our ultra-low-cost design offers a viable option for enhancing mmWave coverage.

5 EVALUATION

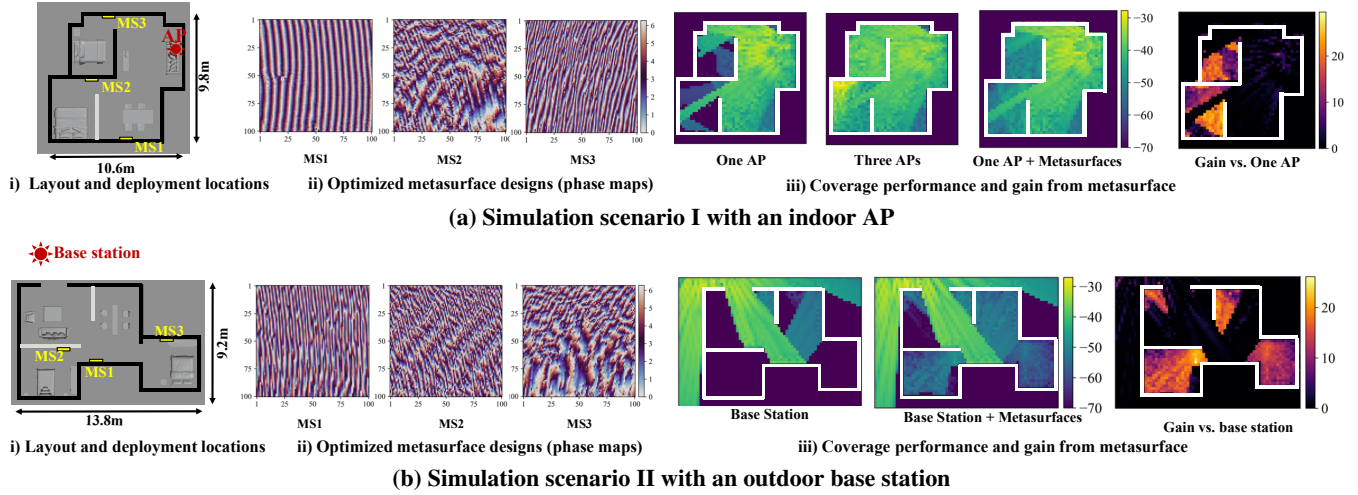


Figure 7: Coverage performance results of our proposed AUTOMS on the simulation models




Figure 8: Channel simulator accuracy




Figure 9: Capacity optimization convergence




Figure 10: Metasurface beam steering validation




Figure 11: Impact of AP antenna array and surface sizes

5.1 Experiment Setup

We use simulation to study diverse environments and settings, and use testbed to evaluate coverage improvement in realistic scenarios. Additionally, we also evaluate AUTOMS’s ability to handle channel fluctuations and compare its performance with alternatives like mirror reflectors and manual placement.

Simulation models. We obtain 3D models of various rooms from a public dataset [18] for AUTOMS analysis. Five models are tested, with the initial two depicted in Figure 7, representing common mmWave coverage solutions: indoor APs and outdoor base stations. In simulations, we use a 15×15 phased array, with a fixed transmitter location and AUTOMS generating the metasurface design, placement, and AP’s codebook.

Testbed setup. For our testbed experiments, we employ TP-Link Talon AD7200 commercial routers [54] equipped with QCA6335 BM [42] as mmWave endpoints (AP and client). By default, we designate the client as the receiver (RX), as metasurface coverage enhancement applies to both uplink and downlink due to channel reciprocity. We flash a modified image of OpenWrt [34, 50, 51] to control routers. For the transmitter, we replace the original antenna array with a 6×6 phased array [60] for better beam steering. By updating the codebook file loaded to the firmware, we change the code words used at AP [49, 55]. Each codeword alters the phase of signals from the phased array antennas to generate

a certain beam pattern. We use 30 codewords for each experiment. The AP selects the best code according to 802.11ad protocol. For the receiver, we use a single antenna with an omni-directional pattern, which matches the default setting of the router. The router driver reports the RSS values of the best codeword to the Linux kernel, which we can access by executing the command `'iw dev $interface scan'` at the RX. This RSS value represents the received signal strength at RX for the beacon from TX. We also set up a separate 2.4 GHz network for router control and remote data collection.

Testbed experiments. We conduct experiments in three indoor scenarios (Figure 14): (1) the AP is in an open office area and AUTOMS extends the coverage to nearby conference rooms, (2) the AP is in a conference room and AUTOMS extends the coverage to the hallway and neighboring room, (3) the AP is in the living room of an apartment and AUTOMS extends the coverage to the bedroom and restroom. We use the ploycam app [39] to obtain the 3D models of the experiment environment. Based on the 3D models, AUTOMS generates the optimal metasurface phase map and deployment locations. We fabricate and deploy the metasurfaces according to our optimization results. Since our commercial APs can only report RSS above -70 dBm, we use -70 dBm as the baseline value for RSS gain calculation when the client fails to report RSS, *i.e.*, blind spots.




Figure 12: Coverage gain comparison between optimized and manually selected placement of metasurfaces

5.2 Micro benchmarks

Channel modeling accuracy. We evaluate the simulator accuracy based on the RSS error between real-world measurement and simulator output. In the real-world scenario III, we test with measurements from 133 locations. Initially, our simulator yields a median error of 4.10 dB using the reflection coefficients from the literature. After estimating these coefficients via gradient descent based optimization (subsection 3.3), we achieve a reduced median error of 2.37 dB in the testing data as shown in Figure 8. Our accuracy aligns with state-of-the-art approach (e.g., Nerf2 [61]), which reports 2.6 dB error.

Optimization convergence time. Figure 9 shows the optimization results over time for three real-world scenarios. We normalize the optimization objective metric, *i.e.*, capacity, and observe that all three scenarios converge within 8 hours. The scatter points show the results after each iteration of the hyper-configuration tuner, which runs the simulated annealing algorithm. Each iteration takes about 1 min to simulate the channel response matrices and compute Equation 1. The total time increases with the number of annealing iterations, but can potentially be reduced significantly through computation reuse across iterations.

Metasurface design validation. To validate our metasurface design, we conduct beam steering experiments. We use two metasurfaces: one that reflects the beam from the AP at 45 degrees and another that produces a wide reflection beam from 30 to 60 degrees. Figure 10 shows that the experimental angles match the target angles well. We compare them with a metal plate that reflects the beam at 0 degrees. Since we measure the results in a conference room instead of an anechoic chamber, we do not report the exact gain of the metasurface beam steering. We only demonstrate the beam steering directions. Our results confirm that our metasurface design can apply the desired phase shifts to the reflected signals.

Impact of phased array and metasurface sizes. In simulation scenario I, we analyze coverage performance gains, *i.e.*, RSS gains, with three metasurfaces at optimal locations, varying the sizes of the phased array and the metasurfaces. Assuming the optimal locations remain consistent across different setups, our results, presented in Figure 11, indicate that gains improve with the enlargement of both the AP phased array and metasurface sizes. With a smaller phased array, a




Figure 13: Coverage and gain comparison between large reflectors and metasurfaces, utilizing median room RSS as a performance metric

metasurface measuring 150×150 nearly maximizes performance, with marginal benefits from further size increases. Conversely, when the metasurface size is relatively large, the performance of coverage can be improved by increasing the phased array size.

Comparison to large mirror reflectors. We compare metasurfaces with large mirror reflectors for coverage performance using simulation scenario I, accounting for a vertical offset between the AP and the target area. The comparison results are shown in Figure 13, we find that metasurfaces provide superior coverage, with RSS gains 7 dB higher than mirror reflectors four times their size. Even large reflectors ($2 \times 2 \text{ m}^2$) are unable to address all blind spots. In contrast to mirror reflectors which are limited to signal reflection at angles dictated by the law of reflection, metasurfaces offer the flexibility to tailor 3D beam directions.

Comparison to manual placement. To assess the impact of metasurface placement on coverage and validate our system’s efficacy, we conduct experiments in simulation scenario II. We select two distinct placement settings for each of the three metasurfaces based on expert knowledge. These placements are then benchmarked against those optimized by our system, as depicted in Figure 12. The findings suggest that the choice of placement caps the metasurface’s potential gains. With our system’s optimization of metasurface and AP phase configurations, we attain median gains of 6.7 dB and 8.9 dB for manual placements 1 and 2, respectively, and 11.6 dB for the optimized placement, all with identical metasurface sizes, highlighting our system’s advantage over manual methods.

5.3 Enhancing mmWave Coverage

Simulated coverage improvement. Figure 7 displays the coverage performance, *i.e.*, RSS distribution heatmaps, with and without metasurfaces for two simulation scenarios. Metasurfaces, when optimized and deployed via AUTOMS, create a coverage comparable to multiple APs but at a significantly reduced cost, virtually eliminating blind spots. Compared to a single AP deployment, AUTOMS achieves substantial RSS

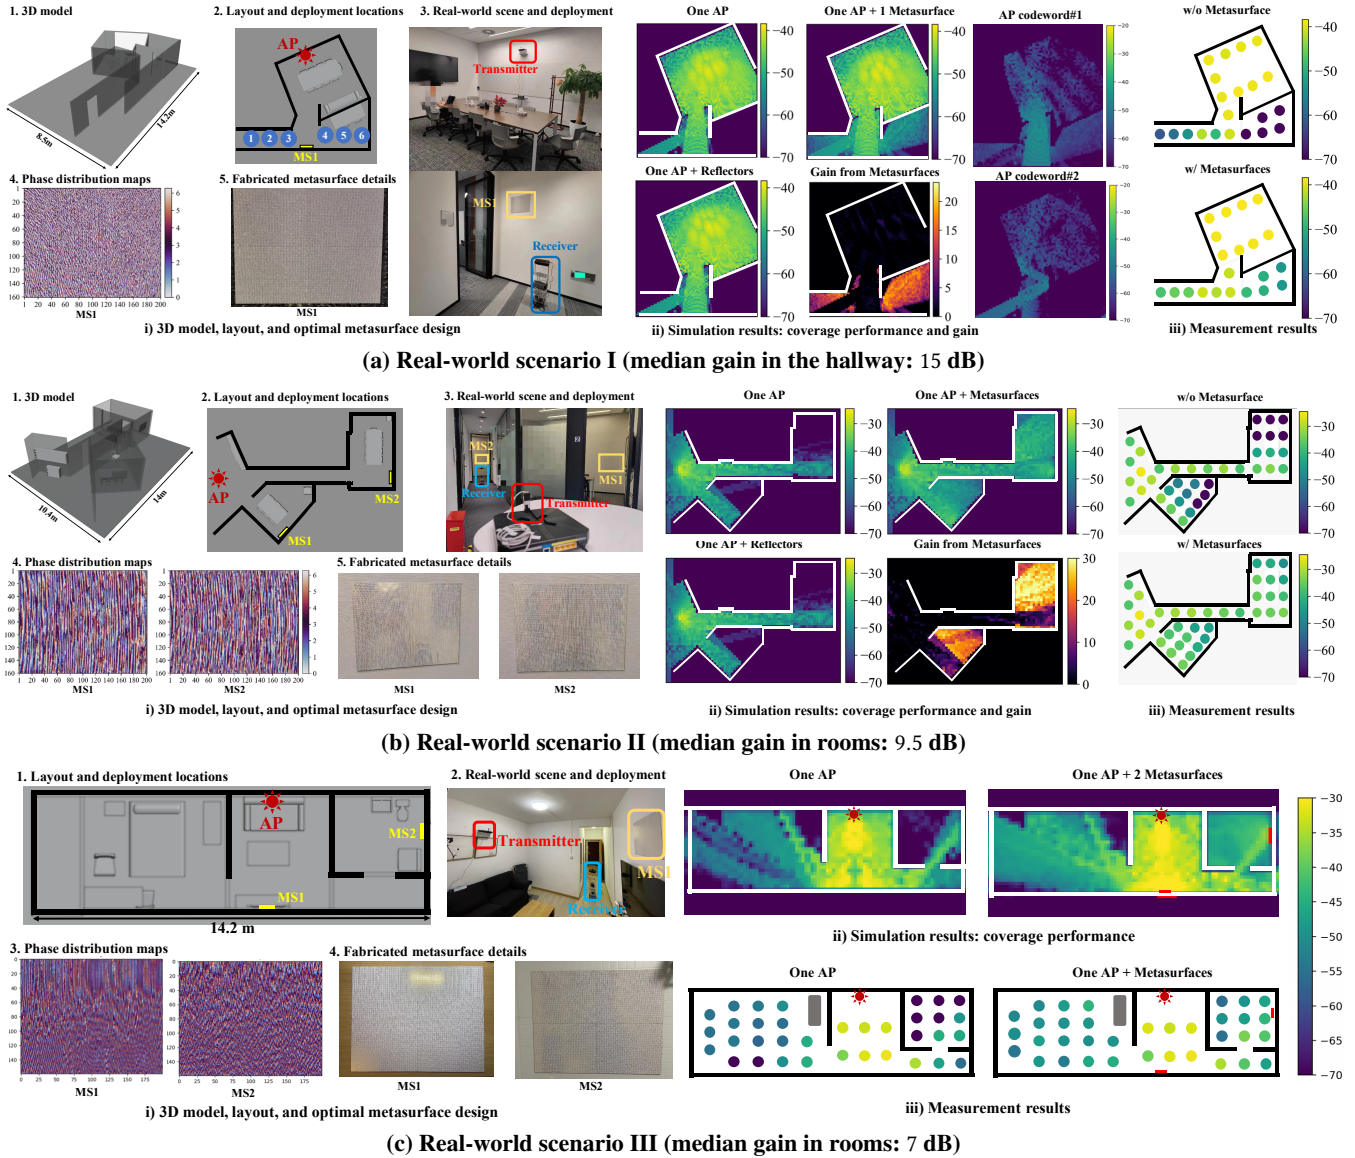


Figure 14: Coverage performance results of our proposed AUTOMS in the real-world scenarios.

gains in target rooms, *i.e.*, rooms not fully covered by LOS paths from a single AP. As evidenced in Figure 15, across five simulation scenarios, we enhance the median RSS by 21.6 dB in target rooms, with a median RSS gain of 12.1 dB, by deploying metasurfaces.

Testbed coverage improvement. In Figure 14, we present our experimental setups, fabricated metasurfaces, and both simulated and measured RSS data with and without metasurfaces for three real-world scenarios. The baseline RSS values with the AP alone align closely with our simulation results for all scenarios. In areas with blocked LOS paths, baseline RSS is low, leading to connectivity issues, so we focus on such areas. Deployment of metasurfaces, however, significantly

improves mmWave coverage in such *target areas*-hallways (real-world scenario I), adjacent conference rooms (real-world scenario II), and neighboring rooms (real-world scenario III), corroborating our simulations. As Figure 15 illustrates, AUTOMS increases the median RSS by 11 dB in these target areas, and achieves gains exceeding 20 dB at previously blind spots compared to the baseline without metasurfaces.

Multi-channel support. We collect RSS measurements on both channel 1 (57.24-59.40 GHz) and channel 2 (59.40-61.56 GHz) of the 802.11ad protocol in real-world scenario I. Our surface design’s wideband phase control yields comparable RSS gains for both channels, with a median increase of 9.5 dB and peaks up to 30 dB, as displayed in Figure 16.




Figure 15: RSS gain for 5 simulation models and 3 real-world scenarios




Figure 16: Gain for two separate mmWave channels




Figure 17: TCP Throughput improvement




Figure 18: RSS measurements over time when people walk around




Figure 19: RSS stability at multiple locations when people walk around




Figure 20: RSS values when the door opens at different angles

Throughput improvement. Next, we show the throughput improvement brought by AUTOMS. We run `iperf3` TCP measurements for RX positions inside conference rooms from scenario II. We notice that the commercial mmWave devices fail to establish a network or TCP connection when RSS is lower than -55 dBm. Our system benefits most to the scenarios where there are low or no throughput, and it can improve the RSS to provide a reliable TCP connection with high throughput. As shown in Figure 17, AUTOMS improves median TCP throughput from 77.2 Mbps without metasurfaces to 373 Mbps with metasurfaces.

5.4 Robustness to Dynamics

In this part, we evaluate the robustness of AUTOMS when facing environmental changes.

Human movements. In testbed scenario I, following optimization outputs, we install the AP and metasurface over 2 m high, above areas of human traffic, enabling signal reach to the metasurface. The metasurface then redirects the beam downward to cover RX locations. We perform 100 RSS measurements at 6 hallway positions, as marked in Figure 14a, at a rate of one per second amidst pedestrian traffic. Figure 18 illustrates the temporal RSS stability at location 1 with the metasurface in place versus without. AUTOMS maintains higher and more consistent RSS levels over time than the baseline. In a static environment, these curves would appear flat. Figure 19 confirms AUTOMS's consistent performance enhancement across multiple locations with minimal RSS fluctuation, even with human movement.

Door Movements. We evaluate the performance of AUTOMS in simulation scenario I, featuring 3 doors at varying angles. Simulations afford exact control over door angles, enabling extensive testing. Findings depicted in Figure 20 reveal that

baseline RSS is affected by door positions, whereas AUTOMS delivers stable and consistently superior RSS compared to the baseline.

6 CONCLUSION

Optimizing mmWave deployment for comprehensive coverage is a significant hurdle for the broad implementation of 5G networks. These challenges stem from the inherently complex nature of mmWave signal propagation as well as the specialized expertise required for effective network management. This paper presents AUTOMS, an innovative framework that automates optimal network deployment using ultra-low-cost passive metasurfaces to enhance mmWave coverage. AUTOMS significantly improves network performance, eliminates coverage dead zones, provides high throughput, and adapts to diverse channel conditions in various settings. Thus, AUTOMS emerges as a promising solution to accelerate the deployment and adoption of mmWave networks.

ACKNOWLEDGMENTS

We are grateful to anonymous reviewers for their constructive comments. We also thank Dr. Yifei Shen for his enlightening discussions.

REFERENCES

- [1] Next G Alliance Report: Roadmap to 6G. <https://nextgalliance.org/wp-content/uploads/2022/02/NextGA-Roadmap.pdf>.
- [2] Omid Abari, Dinesh Bharadia, Austin Duffield, and Dina Katabi. Enabling High-Quality Untethered Virtual Reality. In *NSDI*, pages 531–544, 2017.
- [3] Altair. Altair's 2022.2 Simulation Software. <https://altair.com/simulation-2022-2>, 2023.
- [4] Amazon. NETGEAR Nighthawk X10 Smart WiFi Router (R9000). <https://www.amazon.com/NETGEAR-Nighthawk-Quad>

- Stream-Quad-core-Compatible/dp/B01M12RE4A, 2023.
- [5] Venkat Arun and Hari Balakrishnan. RFocus: Beamforming Using Thousands of Passive Antennas. In *NSDI*, pages 1047–1061, 2020.
 - [6] GSM Association. 5G mmWave Coverage Solutions. https://www.gsm.com/futurenetworks/wp-content/uploads/2022/12/5gmmWave-coverage-extension-solutions-whitepaper_FINAL-v1.pdf, 2023.
 - [7] Ertugrul Basar and Ibrahim Yildirim. Reconfigurable intelligent surfaces for future wireless networks: A channel modeling perspective. *IEEE Wireless Communications*, 28(3):108–114, 2021.
 - [8] Justin Chan, Changxi Zheng, and Xia Zhou. 3D Printing Your Wireless Coverage. In *Proc. of HotWireless*, 2015.
 - [9] Lili Chen, Wenjun Hu, Kyle Jamieson, Xiaojiang Chen, Dingyi Fang, and Jeremy Gummeson. Pushing the physical limits of IoT devices with programmable metasurfaces. In *Symposium on Networked Systems Design and Implementation (NSDI)*, pages 425–438. USENIX, 2021.
 - [10] Lili Chen, Bozhong Yu, Ju Ren, Jeremy Gummeson, and Yaoxue Zhang. Towards Seamless Wireless Link Connection. In *Proceedings of the 21st Annual International Conference on Mobile Systems, Applications and Services*, pages 137–149, 2023.
 - [11] Kun Woo Cho, Mohammad H Mazaheri, Jeremy Gummeson, Omid Abari, and Kyle Jamieson. mmWall: A steerable, transmissive metamaterial surface for NextG mmWave networks. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 1647–1665, 2023.
 - [12] Apple Developer. RoomPlan API. <https://developer.apple.com/augmented-reality/roomplan/>, 2023.
 - [13] Manideep Dunna, Chi Zhang, Daniel Sievenpiper, and Dinesh Bharadia. ScatterMIMO: Enabling virtual MIMO with smart surfaces. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, pages 1–14, 2020.
 - [14] Esteban Egea-Lopez, Jose Maria Molina-Garcia-Pardo, Martine Lienard, and Pierre Degauque. Opal: An open source ray-tracing propagation simulator for electromagnetic characterization. *Plos one*, 16(11):e0260060, 2021.
 - [15] Chao Feng, Xinyi Li, Yangfan Zhang, Xiaojing Wang, Liqiong Chang, Fuwei Wang, Xinyu Zhang, and Xiaojiang Chen. RFlens: metasurface-enabled beamforming for IoT communication and sensing. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*. ACM, 2021.
 - [16] Forsk. Aster Propagation Model. <https://www.forsk.com/aster-propagation-model>, 2023.
 - [17] Forsk. CrossWave Propagation Model. <https://www.forsk.com/crosswave-propagation-model>, 2023.
 - [18] Huan Fu, Bowen Cai, Lin Gao, Ling-Xiao Zhang, Jiaming Wang, Cao Li, Qixun Zeng, Chengyue Sun, Rongfei Jia, Binqiang Zhao, et al. 3d-front: 3d furnished rooms with layouts and semantics. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10933–10942, 2021.
 - [19] Yasaman Ghasempour, Claudio RCM Da Silva, Carlos Cordeiro, and Edward W Knightly. IEEE 802.11 ay: Next-generation 60 GHz communication for 100 Gb/s Wi-Fi. *IEEE Communications Magazine*, 55(12):186–192, 2017.
 - [20] Jean-Baptiste Gros, Vladislav Popov, Mikhail A Odit, Vladimir Lenets, and Geoffroy Lerosey. A reconfigurable intelligent surface at mmwave based on a binary phase tunable metasurface. *IEEE Open Journal of the Communications Society*, 2:1055–1064, 2021.
 - [21] Hichem Guerboukha, Yasith Amarasinghe, Rabi Shrestha, Angela Pizuto, and Daniel M. Mittleman. High-volume rapid prototyping technique for terahertz metallic metasurfaces. *Optics Express*, 29(9):13806–13814, April 2021.
 - [22] Ish Kumar Jain, Raghav Subbaraman, and Dinesh Bharadia. Two beams are better than one: Towards reliable and high throughput mmWave links. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, pages 488–502, 2021.
 - [23] Wahab Khawaja, Ozgur Ozdemir, Yavuz Yapici, Fatih Erden, and Ismail Guvenc. Coverage enhancement for NLOS mmWave links using passive reflectors. *IEEE Open Journal of the Communications Society*, 1:263–281, 2020.
 - [24] B Langen, G Lober, and W+ Herzig. Reflection and transmission behaviour of building materials at 60 ghz. In *5th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, Wireless Networks-Catching the Mobile Future.*, volume 2, pages 505–509. IEEE, 1994.
 - [25] Wonwoo Lee, Semin Jo, Kanghyeok Lee, Hong Soo Park, Junhyuk Yang, Ha Young Hong, Changkun Park, Sun K Hong, and Hojin Lee. Single-layer phase gradient mmwave metasurface for incident angle independent focusing. *Scientific reports*, 11(1):12671, 2021.
 - [26] Zhuqi Li, Yaxiong Xie, Longfei Shangguan, Rotman Ivan Zelaya, Jeremy Gummeson, Wenjun Hu, and Kyle Jamieson. Towards programming the radio environment with large arrays of inexpensive antennas. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*, pages 285–300, 2019.
 - [27] Christos Liaskos, Shuai Nie, Ageliki Tsioliaridou, Andreas Pitsillides, Sotiris Ioannidis, and Ian Akyildiz. End-to-end wireless path deployment with intelligent surfaces using interpretable neural networks. *IEEE Transactions on Communications*, 68(11):6792–6806, 2020.
 - [28] Christos Liaskos, Ageliki Tsioliaridou, Shuai Nie, Andreas Pitsillides, Sotiris Ioannidis, and Ian F Akyildiz. On the network-layer modeling and configuration of programmable wireless environments. *IEEE/ACM transactions on networking*, 27(4):1696–1713, 2019.
 - [29] Yuchen Liu, Yubing Jian, Raghupathy Sivakumar, and Douglas M Blough. Maximizing line-of-sight coverage for mmWave wireless LANs with multiple access points. *IEEE/ACM Transactions on Networking*, 30(2):698–716, 2021.
 - [30] Ruichun Ma and Wenjun Hu. RF-Mediator: Tuning Medium Interfaces with Flexible Metasurfaces. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, 2024.
 - [31] Ruichun Ma, R. Ivan Zelaya, and Wenjun Hu. Softly, Deftly, Scrolls Unfurl Their Splendor: Rolling Flexible Surfaces for Wideband Wireless. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, 2023.
 - [32] Ajay Mahimkar, Ashwan Sivakumar, Zihui Ge, Shomik Pathak, and Karunasish Biswas. Auric: using data-driven recommendation to automatically generate cellular configuration. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, pages 807–820, 2021.
 - [33] Microsoft Research. Neural Network Intelligence. <https://github.com/microsoft/nni>.
 - [34] Joan Palacios, Daniel Steinmetzer, Adrian Loch, Matthias Hollick, and Joerg Widmer. Adaptive codebook optimization for beam training on off-the-shelf ieee.11 ad devices. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, pages 241–255, 2018.
 - [35] Hao Pan, Lili Qiu, Bei Ouyang, Shicheng Zheng, Yongzhao Zhang, Yi-Chao Chen, and Guangtao Xue. PMSat: Optimizing Passive Metasurface for Low Earth Orbit Satellite Communication. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, pages 1–15, 2023.
 - [36] Steven G Parker, James Bigler, Andreas Dietrich, Heiko Friedrich, Jared Hoberock, David Luebke, David McAllister, Morgan McGuire, Keith Morley, Austin Robison, et al. Optix: a general purpose ray tracing engine. *Acmm transactions on graphics (tog)*, 29(4):1–13, 2010.
 - [37] Zhangyou Peng, Linxiao Li, Miao Wang, Zhonghao Zhang, Qi Liu, Yang Liu, and Ruoran Liu. An effective coverage scheme with passive-reflectors for urban millimeter-wave communication. *IEEE Antennas*

- and *Wireless Propagation Letters*, 15:398–401, 2015.
- [38] Carl Pfeiffer and Anthony Grbic. Metamaterial Huygens' Surfaces: Tailoring Wave Fronts with Reflectionless Sheets. *Physical Review Letters*, 110(19):197401, May 2013.
- [39] polycam. Polycam - LiDAR & 3D Scanner for iPhone & Android. <https://poly.cam/>, 2023.
- [40] David M Pozar. *Microwave engineering*. John Wiley & sons, 2011.
- [41] Kun Qian, Lulu Yao, Xinyu Zhang, and Tse Nga Ng. MilliMirror: 3D printed reflecting surface for millimeter-wave coverage expansion. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking*, pages 15–28, 2022.
- [42] Qualcomm. 802.11ad MAC baseband transceiver. <https://www.qualcomm.com/products/technology/wi-fi/qca6335>, 2023.
- [43] Remcom. Engineered Electromagnetic Surfaces (EES). <https://www.remcom.com/wireless-insite-engineered-electromagnetic-surfaces-ees>, 2023.
- [44] Remcom. Wireless InSite 3D Wireless Prediction Software. <https://www.remcom.com/wireless-insite-em-propagation-software/>, 2023.
- [45] Katsuyoshi Sato, Takeshi Manabe, Toshio Ihara, Hiroshi Saito, Shigeru Ito, Tetsu Tanaka, Kazuyoshi Sugai, Norichika Ohmi, Yasushi Murakami, Masanori Shibayama, et al. Measurements of reflection and transmission characteristics of interior structures of office building in the 60-ghz band. *IEEE transactions on antennas and propagation*, 45(12):1783–1792, 1997.
- [46] Sivers Semiconductors. Sivers Evaluation Kits (EVK) and Evaluation Boards (EVB). <https://www.sivers-semiconductors.com/sivers-wireless/evaluation-kits/>, 2023.
- [47] Zhambyl Shaikhanov, Sherif Badran, Hichem Guerboukha, Josep Jornet, Daniel Mittleman, and Edward Knightly. MetaFly: Wireless Backhaul Interception via Aerial Wavefront Manipulation. In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 151–151. IEEE Computer Society, 2024.
- [48] Zhambyl Shaikhanov, Fahid Hassan, Hichem Guerboukha, Daniel Mittleman, and Edward Knightly. Metasurface-in-the-Middle Attack: From Theory to Experiment. In *Proceedings of the 15th ACM Conference on Security and Privacy in Wireless and Mobile Networks, WiSec '22*, pages 257–267, New York, NY, USA, May 2022. Association for Computing Machinery.
- [49] Yiwen Song, Changhan Ge, Lili Qiu, and Yin Zhang. 2ACE: Spectral Profile-driven Multi-resolutional Compressive Sensing for mmWave Channel Estimation. In *Proceedings of the Twenty-fourth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*, pages 41–50, 2023.
- [50] Daniel Steinmetzer, Daniel Wegemer, and Matthias Hollick. Talon tools: The framework for practical IEEE 802.11ad research, 2018.
- [51] Daniel Steinmetzer, Daniel Wegemer, Matthias Schulz, Joerg Widmer, and Matthias Hollick. Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices. In *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies, CoNEXT '17*, pages 414–425, New York, NY, USA, November 2017. Association for Computing Machinery.
- [52] Mushtaq Talib, Norazizah Binti Mohd Aripin, Noor Shamsiah Othman, and Adheed Hasan Sallomi. Comprehensive overview on millimeter wave communications for 5g networks concentrating on propagation models for different urban environments. In *Journal of Physics: Conference Series*, volume 2322, page 012095. IOP Publishing, 2022.
- [53] Xin Tan, Zhi Sun, Dimitrios Koutsonikolas, and Josep M Jornet. Enabling indoor mobile millimeter-wave networks based on smart reflect-arrays. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*, pages 270–278. IEEE, 2018.
- [54] Tp-link. Talon AD7200 Multi-Band Wi-Fi Router. <https://www.tp-link.com/sg/home-networking/wifi-router/ad7200/>, 2023.
- [55] Teng Wei and Xinyu 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*, pages 42–55, 2017.
- [56] Teng Wei, Anfu Zhou, and Xinyu Zhang. Facilitating Robust 60 GHz Network Deployment By Sensing Ambient Reflectors. In *NSDI*, pages 213–226, 2017.
- [57] Allen Welkie, Longfei Shangguan, Jeremy Gummesson, Wenjun Hu, and Kyle Jamieson. Programmable Radio Environments for Smart Spaces. In *Workshop on Hot Topics in Networks (HotNets)*. ACM, 2017.
- [58] R Ivan Zelaya, Ruichun Ma, and Wenjun Hu. Towards 6G and Beyond: Smarten Everything with Metamorphic Surfaces. In *Proceedings of the Twentieth ACM Workshop on Hot Topics in Networks*, pages 155–162, 2021.
- [59] R Ivan Zelaya, William Sussman, Jeremy Gummesson, Kyle Jamieson, and Wenjun Hu. LAVA: fine-grained 3D indoor wireless coverage for small IoT devices. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, pages 123–136, 2021.
- [60] Renjie Zhao, Timothy Woodford, Teng Wei, Kun Qian, and Xinyu Zhang. M-cube: A millimeter-wave massive MIMO software radio. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, pages 1–14, 2020.
- [61] Xiaopeng Zhao, Zhenlin An, Qingrui Pan, and Lei Yang. NeRF2: Neural Radio-Frequency Radiance Fields. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, 2023.
- [62] Xia Zhou, Zengbin Zhang, Yibo Zhu, Yubo Li, Saipriya Kumar, Amin Vahdat, Ben Y Zhao, and Haitao Zheng. Mirror mirror on the ceiling: Flexible wireless links for data centers. *ACM SIGCOMM Computer Communication Review*, 42(4):443–454, 2012.