FlowForm: Scalable Passive Metasurface Network for mmWave Coverage Expansion

Wuqiong Zhao
University of California San Diego
La Jolla, California, USA
wqzhao@ucsd.edu

Baicheng Chen
University of California San Diego
La Jolla, California, USA
b3chen@ucsd.edu

Kai Zheng
University of California San Diego
La Jolla, California, USA
kazheng@ucsd.edu

Xingyu Chen
University of California San Diego
La Jolla, California, USA
xic063@ucsd.edu

Wenqian Zhang
University of California San Diego
La Jolla, California, USA
wez040@ucsd.edu

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

Abstract

Millimeter wave (mmWave) networks offer multi-gigabit data rates but suffer from severe path loss and blockage, resulting in spotty coverage. Emerging reconfigurable intelligent surfaces (RIS) can mitigate these challenges, but their reliance on active control channels, power sources, and complex runtime coordination imposes significant hardware and deployment overhead. This paper introduces FlowForm, a system that expands mmWave coverage using networks of passive metasurfaces that require no power, control, or runtime coordination. FlowForm's key innovation is a hierarchical flow topology that organizes passive metasurfaces into major flows (directional relay chains using near-field focusing) and minor flows (wide-area fan beams), enabling multi-hop passive routing and over-the-air combination of analog signals. We develop a theoretical framework establishing the optimality of this topology and a hierarchical optimization algorithm that jointly determines metasurface placement and beam configurations. FlowForm operates transparently with standard mmWave network protocols, managing channel dynamics and multi-user interference through diversity-aware design rather than runtime reconfiguration. Our experimental evaluation across five indoor environments demonstrates up to 94% average rate improvement and 114% coverage expansion using low-cost 3D-printed metasurfaces ($2 per unit), achieving performance comparable to active RIS at orders of magnitude lower cost.

CCS Concepts

• **Networks → Physical links**; • **Hardware → Networking hardware**; **Wireless devices**.

Keywords

Intelligent reflecting surface, reconfigurable intelligent surface, metasurfaces, mmWave networking, passive relay network

SIGCOMM '26, Denver, CO, USA

ACM ISBN 979-8-4007-2467-1/26/08
https://doi.org/10.1145/3789240.3829102

ACM Reference Format:
Wuqiong Zhao, Baicheng Chen, Kai Zheng, Xingyu Chen, Wenqian Zhang, and Xinyu Zhang. 2026. FlowForm: Scalable Passive Metasurface Network for mmWave Coverage Expansion. In *ACM SIGCOMM 2026 Conference (SIGCOMM '26), August 17–21, 2026, Denver, CO, USA.* ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3789240.3829102

1 Introduction

Millimeter wave (mmWave) communication is fundamental to 5G and beyond, offering multi-gigabit data rates through wide available bandwidth [56]. However, the high-frequency mmWave signals suffer from severe path loss and susceptibility to blockages [27, 55]. While directional beamforming with large phased array antennas can extend communication range [67, 68], the fundamental problem of spotty coverage persists—slight movements or obstructions cause dramatic signal degradation. Traditional solutions like infrastructure densification or active relays remain difficult to scale, as they incur high deployment costs and require complex, low-latency control channels to manage dynamic beam steering.

Metasurfaces, consisting of massive arrays of subwavelength electromagnetic elements, present a promising alternative [54]. By strategically attaching metasurfaces to ambient surfaces like walls and ceilings, signals can be redirected toward previously unreachable blind spots, effectively "programming the wireless environment" [36]. Active metasurfaces, also known as reconfigurable intelligent surfaces (RIS), can adapt their reflection beam patterns in real-time [22, 31]. However, RIS require control circuits, power sources, and sophisticated coordination protocols [31]. In contrast, passive metasurfaces eliminate such needs, and can be fabricated through low-cost 3D-printing or hot-stamping [41, 53]. However, prior work primarily focused on standalone metasurfaces used in isolation. It remains an open question whether such passive metasurfaces can provide robust coverage in moderate to large-scale deployment, especially considering the fundamental limitation of fixed reflecting beams determined at design time.

In this paper, we introduce FlowForm, a system that creates scalable passive metasurface networks to provide low-cost and transparent coverage expansion. Our key research question is: *How can we effectively expand mmWave network coverage through a coordinated network of passive metasurfaces that efficiently propagate signals between each other toward target regions?* This question breaks down into three key challenges.

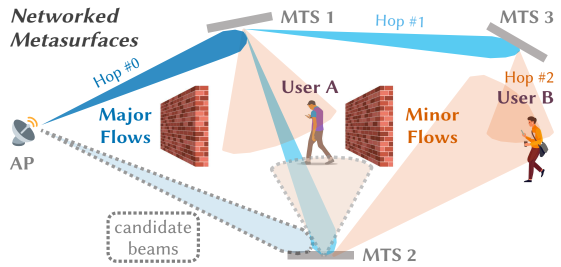

Figure 1: FlowForm passive metasurface (MTS) network illustration: Major flows dominate spatial coverage, minor flows fill in orientational blockage, users with merged flow experience boosted throughput.

(i) Network topology design for passive metasurfaces. The passive metasurface network requires a principled topology design that can easily scale while remaining tractable to optimize. Traditional relay network topologies do not account for the analog signal degradation or combination along a chain of passive reflectors. Moreover, passive metasurfaces have fixed reflection beam patterns determined at design time, making classical adaptive routing protocols inapplicable.

To overcome this challenge, FlowForm organizes the metasurfaces into a hierarchical flow topology. The key insight is that passive metasurfaces face a fundamental gain–coverage tradeoff: long-range relay requires narrow, high-gain beams, whereas broad local coverage requires wide fan beams with inherently lower directional gain. Analogous to how tributaries merge into rivers, FlowForm creates communication flows at two scales (Fig. 1). Major flows are highly directional relay paths serving as the communication backbone. They pass the analog signals along a chain of reflection surfaces, while using near-field beam focusing to overcome path loss. Minor flows are wide-angle fan beams from metasurfaces directly to target receiver locations, filling in coverage gaps while providing diverse signal paths towards target locations. We develop a theoretical framework that establishes the optimality of the major-minor flow topology, which emerges from the physics of passive relaying and power splitting in metasurface networks.

(ii) Scalable optimization of metasurface placement and beam configuration. For a given deployment site, we must jointly determine the metasurface placement and unit-cell configurations (thousands of elements per metasurface), which creates a mixed discrete-continuous problem with exponential complexity.

We address this challenge through two key techniques. First, we develop a metasurface abstraction framework that reduces dimensionality by using a compact parametric form. Second, we design a hierarchical optimization framework to decompose the problem into sequential stages: major flow design identifies critical relay bottlenecks through topology analysis and near-field beam optimization; then minor flow design completes coverage through a novel annealed passive routing optimization algorithm. This hierarchical approach reduces complexity from exponential to approximately linear in the number of metasurfaces.

(iii) Handling channel dynamics and interference without metasurface reconfiguration. How can a static metasurface network provide robust coverage under receiver mobility, time-varying fading, and multi-user interference?

We answer this question by designing for diversity rather than adaptation, and by leveraging the active radios' beam steering capability. With this design, the metasurface network can operate transparently to standard mmWave radios, without modification to the hardware or network protocols. Our optimization framework explicitly creates diverse signal arrival angles towards target locations. Standard beam scanning procedures at the transceivers dynamically select among these pre-engineered paths to ensure robustness. For multi-user scenarios, we introduce interference-aware optimization that preconfigures metasurface patterns to favor spatial sharing.

We implement FlowForm using low-cost 3D-printed metasurfaces and validate it through extensive experiments in 5 diverse environments. FlowForm achieves up to 94% average link capacity improvement and 114% network coverage expansion compared to state-of-the-art baselines. Compared to active RIS, FlowForm achieves comparable coverage at orders of magnitude lower cost. FlowForm maintains stable coverage as user devices move and rotate, which also indirectly reduces the beam scanning overhead of the access point (AP). Moreover, by optimizing for spatial reuse opportunities, FlowForm achieves an additional 43% throughput gain in a multi-AP multi-user network, given the same coverage area. The system tolerates environment modeling errors and placement inaccuracies, thus enabling simple and practical deployment.

This work makes the following contributions:

(i) We introduce the major–minor flow topology for passive metasurface networks and prove its optimality, providing a principled decomposition that enables scalable optimization.

(ii) We propose metasurface abstraction and hierarchical optimization techniques, including annealed passive routing optimization (APRO) that jointly optimize the placement and design-time configuration of the metasurface network for maximum coverage.

(iii) Through a path diversity-aware design, we demonstrate FlowForm's robustness to network dynamics and seamless integration with standard protocols.

(iv) We implement FlowForm and validate its performance through extensive real-world deployments.

2 Related Work

Active relays have been widely studied for wireless coverage expansion [2, 6, 7, 21, 37, 49, 66]. At mmWave frequencies, relay networks face unique challenges from highly directional links and the need for sophisticated protocols to coordinate the directional relays and base stations [7, 66, 70]. Consequently, active relay deployment remains constrained by cost and complexity, and can hardly scale to larger networks.

Active RIS offer dynamic beam steering by adjusting element phase shifts in real-time [12, 16, 17, 22, 31, 33–35, 39]. However, reconfigurability incurs substantial drawbacks: *(i)* control channels for run-time coordination which requires a separate radio on each RIS, bringing complexity comparable to active relays; *(ii)* power supplies and control circuits that increase deployment complexity; and *(iii)* continuous joint optimization requiring full channel state information. These drawbacks are further exacerbated by the prohibitive cost scaling of active components at higher operating

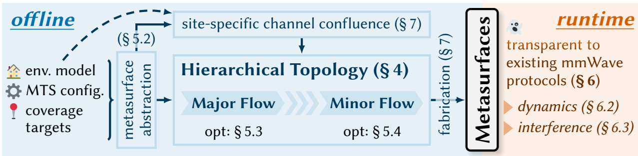

Figure 2: FlowForm system overview.

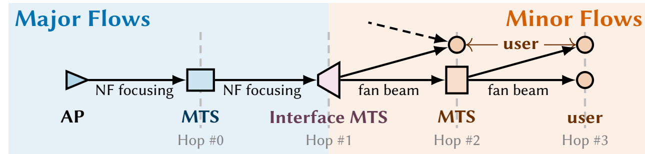

Figure 3: Illustration of FlowForm's proposed major–minor flow topology.

frequencies. RFMagus [34] coordinates multiple RIS to form relay paths, but inherits these fundamental limitations of active RIS.

Passive metasurfaces eliminate these constraints through low-cost fabrication such as 3D-printed dielectrics [53, 65], split-ring resonator (SRR)-based reflectarrays on paper substrates [41, 52], or building material diffraction/reflection [40, 51]. However, prior work focuses on standalone hardware design without considering network-level coordination, channel dynamics or interference. AutoMS [41] represents the closest prior work. It optimizes placement for a single metasurface and beam patterns for up to three metasurfaces. However, this approach does not scale to network-level deployments because independent optimization of metasurfaces prevents coordinated signal reflection and relaying, precluding spatial reuse and interference-aware design for multiple users.

FlowForm combines the best of both worlds: low fabrication cost, minimal deployment complexity, and zero runtime overhead of passive metasurfaces with network-level performance comparable to active RIS, as we will demonstrate through experiments.

3 System Overview

FlowForm targets indoor environments, where sophisticated geometries and obstacles create NLoS blockage challenges for mmWave connectivity. The FlowForm passive metasurfaces (MTS) can be strategically attached to walls and ceilings to redirect signals around these blockages. Fig. 2 illustrates FlowForm's major components and workflow.

FlowForm operates in two phases. *During the offline phase*, given a deployment site with known AP locations, candidate metasurface mounting locations, and target user coverage areas, FlowForm determines the optimal metasurface placement and beam pattern configuration for each metasurface. These inputs are standard for commercial wireless network deployment and are typically obtained from architectural floor plans and a brief site survey [1, 23, 28, 30, 50, 63]. The optimization follows three key stages: *(i)* Metasurface abstraction (§ 5.2) reduces the parameter space from thousands of element-wise phase shifts to a compact parametric representation. *(ii)* Major flow design (§ 5.3) establishes directional relay chains through beamforming-aware path planning and near-field focusing optimization. *(iii)* Minor flow design (§ 5.4) completes spatial coverage through annealed passive routing optimization (APRO), which optimizes the metasurfaces to provide wide coverage while minimizing redundancy. Throughout this process, the optimization accounts for practical fabrication constraints (§ 7), including phase quantization and efficiency limitations.

At runtime, FlowForm metasurfaces operate transparently to standard mmWave protocols. From the active radios' perspectives, metasurfaces behave like natural reflectors which simply redirect signals according to preconfigured beam patterns. Users maintain direct communication with APs using standard protocols [24, 25], including beam scanning and MAC coordination. For multi-AP, multi-user scenarios, FlowForm's offline optimization explicitly accounts for interference (§ 6.3), by optimizing the metasurface patterns to encourage spatial sharing opportunities across possible user locations while minimizing interference. Without loss of generality, our design and experimental validation employ the mmWave network protocol defined in IEEE 802.11ad [24] and 802.11ay [25].

4 Topology Design for Networked Metasurfaces

4.1 Can Passive Metasurface Relays Beat the Path Loss?

Passive metasurfaces do not amplify signals—they only redistribute incident energy. Therefore, end-to-end performance is ultimately constrained by free-space path loss. Increasing the metasurface size helps capture more energy from the transmitter, yet it cannot surpass the fundamental performance bound. However, multiple distributed metasurfaces can maintain the link budget from transmitter to receiver by forming coherent multi-hop passive relay paths.

This approach provides three key advantages. *First*, it enables signals to be steered around obstacles and into otherwise unreachable areas. *Second*, multi-hop relaying replaces one long lossy link with several short links, where each metasurface's near-field focusing gain can be effectively exploited to overcome the far-field multiplicative path loss. Suppose d_i is the distance of the i-th link along a K-hop path, near-field focusing can approximately transition the link budget from a product of distances $\propto (d_1 \cdot d_2 \cdots d_K)^2$ to an additive distance loss $\propto (d_1 + d_2 + \cdots + d_K)^2$ [29]. This allows the multi-hop passive relays to mimic the performance of a continuous LoS path. *Third*, as we will show in § 6.3, distributed metasurfaces can improve the path diversity, which is essential to improve robust coverage and spatial multiplexing.

Notably, these benefits arise only when metasurfaces are judiciously designed and deployed. FlowForm explicitly optimizes metasurface placement and reflection beam patterns to create such multihop paths that relay and combine the analog signals over the air.

4.2 Major–Minor Flow Topology

Major flows vs. minor flows. Designing an effective relay network topology for passive metasurfaces presents unique challenges. Traditional relay network designs are typically optimized for active radios, which can amplify and regenerate signals. In contrast, the performance of FlowForm is dominated by cumulative path loss and the collective gain of the coordinated metasurfaces. In addition, traditional network topology design does not account for the unique physics of passive reflections, such as the near-field (NF) focusing. Moreover, passive metasurfaces cannot perform multi-user coordination, requiring explicit offline network topology design to enable spatial reuse through beam steering at active transmitters and receivers.

To address these challenges, we propose the major–minor flow topology, as illustrated in Fig. 3. This topology organizes metasurfaces into two functional categories.

Major flows are highly directional, focused beams established between the access point (AP) and metasurfaces (MTS), or between metasurfaces themselves. They serve as the "backbone" of the network, relaying signals across long distances while exploiting near-field focusing to overcome path loss.

Minor flows are fan beams that provide wide-area coverage from metasurfaces directly to user locations. These flows prioritize angular coverage over directional gain, ensuring that end users can access the network from diverse positions and orientations. They essentially fill in the coverage gaps left by the major flows, particularly in regions beyond the AP's direct reach where major flow relay chains terminate.

Optimality of the major–minor flow topology. We now establish a theoretical model showing that optimal solutions naturally exhibit the major–minor flow structure.

Without loss of generality, we consider the downlink case here. Consider a network with a single AP as transmitter (Tx), a set $\mathcal{M}$ of candidate metasurface locations, and a set $\mathcal{R}$ of receiver (Rx) locations. We model the reflection beam pattern of a metasurface by the relative power of signals along each discretized angle within the metasurface's field of view (FoV). Specifically, denote by $S_i[k] \in [0, 1]$ the relative signal power distributed along angular direction k at metasurface i. Our goal is to maximize the sum rate $\mathcal{U}(\cdot)$ across all Rx locations (with SNR_r for receiver r) subject to the power distribution constraints, *i.e.*,

$$\begin{aligned} \textbf{(P)} \quad & \max_{\{S_i\}} \mathcal{U}(S), \quad \mathcal{U}(S) \triangleq \sum_{r\in\mathcal{R}} \log(1 + \mathrm{SNR}_r(S)), \\ & \text{s.t. } S_i[k] \ge 0, \quad \sum_k S_i[k] \le 1, \quad \forall i \in \{0\} \cup \mathcal{M}. \end{aligned} \tag{1}$$

While **(P)** maximizes aggregate capacity, coverage and fairness across receiver locations are achieved through the diversity scoring mechanism in § 6.2, which weights the objective to encourage multiple paths serving diverse locations and orientations.

We show that optimal power allocations exhibit major–minor flow structure through two steps. *First*, we derive first-order optimality conditions using Gateaux derivatives (Lemma 4.1), which characterize the directional scores that guide optimal power allocation. *Second*, we show that under typical indoor propagation conditions, these directional scores favor the major–minor flow structure (Proposition 4.2). The proofs leverage convex optimization techniques (KKT conditions) and the physical properties of mmWave propagation.

Let $y_i(S)$ denote the outgoing power at i under S. Using Gateaux derivative [38], we have

$$\frac{\partial \mathcal{U}}{\partial S_i[k]} = y_i(S)\Delta_i[k], \quad \Delta_i[k] \triangleq R_i[k] + D_i[k], \tag{2}$$

where $\Delta_i[k]$ is the directional score, which models the sum of direction power projected to other metasurfaces (*i.e.*, $R_i[k]$) and to the receivers (*i.e.*, $D_i[k]$). With this setup, we have:

Lemma 4.1 (First-Order Optimality). *At any local/global optimizer $S^\star$ of **(P)**, if $S_i^\star[k] > 0$ then $\Delta_i^\star[k] = \sup_{k'} \Delta_i^\star[k']$, where $\Delta_i[k]$ is the directional score. (Proof in § B.2).*

This lemma formalizes a waterfilling principle for passive metasurface networks. Optimal power allocation should focus energy on the directions with highest directional scores providing maximum marginal utility. The major–minor flow structure emerges when, for upstream metasurfaces, the relay value $R_i[k]$ dominates the direct coverage $D_i[k]$; while for downstream, the reverse is true. The sufficient condition for major flow behavior at metasurface i is:

$$\max_k \Delta_i^\star[k] > \max_k D_i^\star[k], \tag{3}$$

meaning that the best directional choice involves relaying rather than direct coverage.

In typical indoor geometries, each angular bin k has a dominant first-interaction class: distant metasurfaces have high path loss but subtend larger user regions (receiver-facing bins), while nearby metasurfaces have low path loss but cover fewer users (MTS-facing bins). We formalize this as condition G_0 (angular separability) in Appendix B.3.

Proposition 4.2 (Major–Minor Flow Optimality). *Under G_0, any optimizer of **(P)** allocates power exclusively to either MTS-facing or receiver-facing angular bins at each node, yielding the major–minor flow topology; when G_0 holds approximately (as in typical indoor environments), this structure remains near-optimal. (Proof in § B.3).*

This separation yields the major–minor topology: upstream metasurfaces concentrate power into high-gain relay chains, while downstream nodes fan out for broad coverage, maximizing both reach and diversity. In practice, this separation may not be strict at every node, which our optimization accounts for (§ 5.4).

5 Major-Minor Flow Optimization

5.1 Optimization Overview

Designing a network of passive metasurfaces requires solving a challenging joint optimization problem: *(i)* selecting the location of metasurfaces and *(ii)* determining the reflection beam pattern of each. The beam pattern of a passive metasurface is in turn determined by the phase shift configurations of thousands of unit-cell elements. Jointly optimizing the metasurfaces' locations and the element-wise phase configurations entails a mixed discrete-continuous problem with exponential complexity.

FlowForm addresses this complexity by sparsifying the parameter space, and by leveraging the decomposed major–minor flow structure established in Proposition 4.2.

Optimization parameters. The joint placement and phase optimization operates on the parameters:

$$\mathbf{\Psi} \triangleq [\overbrace{\alpha_1, \dots, \alpha_M}^{\text{placement}}, \overbrace{\mathbf{p}_1, \dots, \mathbf{p}_M}^{\text{patterns}}], \tag{4}$$

where M is the number of candidate locations, $\alpha_m \in \{0, 1\}$ indicates whether a metasurface is placed at location m, and $\mathbf{p}_m$ is the reflection beam pattern of metasurface m.

To reduce the dimensionality of the parameter space, we employ a metasurface abstraction framework (§ 5.2), which represents metasurfaces' configuration parameters in a compact, physically-motivated parameter space rather than optimizing individual phase elements directly.

Optimization objectives and constraints. FlowForm maximizes the average physical rate across all target locations, which inherently balances coverage (sufficient SNR across locations) and capacity (maximizing throughput where coverage exists). Additionally, we optimize for multiple strong paths to diversify the signal incidental angles at each location, thus providing robustness against receiver orientation and user mobility. The underlying channel model is multipath- and beamforming-aware through site-specific ray tracing (§ 7).

The optimization operates under several practical constraints: ***(i)*** Environmental constraints: metasurface placements are restricted to viable mounting surfaces (walls, ceilings) and cannot overlap (Eq. (4)); ***(ii)*** Metasurface constraints: large incident or reflection angles degrade modeling accuracy and reflection efficiency of metasurfaces [45], so we intentionally constrain the reflection beam's pointing angles (§ 5.2); ***(iii)*** Fabrication constraints: phase patterns must be realizable with quantized phase shifts and account for manufacturing tolerances (§ 5.2, § 7).

Optimization workflow. Decoupling the placement and configurations of metasurfaces is sub-optimal and can lead to poor local minima. AutoMS employs gradient-free optimization for placement [41], yet it does not scale well due to high complexity. Alternatively, one may exhaustively search for the placement with best possible patterns, and greedily add more metasurfaces. This forms a strong yet computationally inhibitive and sub-optimal (not scalable) baseline that we call FlowForm Greedy-Exhaustive (FF GrEx).

In contrast, we propose a hierarchical major–minor optimization method which decomposes the problem into two sequential stages: major flow design optimizes critical long relay paths leveraging near-field beam focusing, while minor flow design completes spatial coverage by optimizing fan beams from the major flows' terminating points (referred to as interface metasurfaces). We summarize the optimization workflow in Alg. 1 and provide more details below.

Algorithm 1 FlowForm Optimization Workflow

```
Require: Scene & AP info, candidate MTS positions, budget
Ensure: Placement {α_m}, fabrication-ready patterns {p_m}
Preparation
    Build site-specific channel confluence model ........ § 7
    Abstract MTS patterns as beam arms .................. § 5.2
Major Flow .............................................. § 5.3
    Identify coverage bottlenecks; plan relay paths
    Optimize NF focused beams along relay chains
    Select interface MTSs via greedy maximum coverage
Minor Flow (APRO) ....................................... § 5.4
    Pre-prune candidates; warm-start patterns
    for temperature β in annealing schedule do
        Compute routing weights W[m, r]
        Update selection {α_m} via contribution ranking
        Update patterns {p_m} via gradient descent:
            w/ diversity scoring ........................ § 6.2
            w/ interference-aware objective ............. § 6.3
    Refine selection via swap evaluation
    Joint phase finetuning
Fabrication ............................................. § 7
    Map to element phases; optimize for fabrication
```

5.2 Metasurface Abstraction

Parameterizing reflection beam patterns. Per Lemma 4.1, optimal metasurface patterns concentrate power in a *sparse* set of directions. We exploit this structure by representing each metasurface's reflection beam pattern as a sum of parametric "beam arms," each defined by just four parameters: steering angle, beamwidth (quadratic coefficient), amplitude, and phase. This representation maintains physical realizability, *i.e.*, these beams can be synthesized even though the metasurface elements can only change the phase and not the magnitude of incidental signals. Meanwhile, it reduces the parameter count from thousands (for element-wise phases) to typically 8–16 (for 2–4 beam arms).

Mathematically, each beam arm is characterized by a quadratic phase profile in the angular domain [18], which produces a focused beam at a target direction with controlled beamwidth. The quadratic coefficient controls beam focusing: larger coefficients produce narrower beams suitable for major flows, while smaller coefficients produce wider beams for minor flow coverage.

This representation provides 3 key advantages: *(i)* It dramatically reduces optimization space; *(ii)* The structured representation prevents overfitting to discretized user locations and produces beam patterns that generalize better to actual deployments; *(iii)* The parameters directly correspond to physical beam characteristics (*e.g.*, direction, width) that eventually map to metasurfaces' fabrication parameters.

Handling non-ideal effects. Real metasurfaces exhibit imperfections: phase quantization from limited fabrication resolution, amplitude variations across elements, coupling effects, and grating lobes. We account for these systematically through an efficiency parameter $\eta \in (0, 1]$ that scales the ideal beam pattern, combined with explicit modeling of dominant artifacts (*e.g.*, grating lobes). Specifically, for common passive metasurfaces that form beams through element-wise phase shifts [13], grating lobes appear at angles θ_g satisfying $d(\sin\theta - \sin\theta_g) = \pm n\lambda$, where d is the element spacing, θ is the designed main beam angle, and n is the grating order [8, 61] due to unwanted periodic amplitude variations. For reliability, we do not count the grating lobe for coverage during minor flow optimization; but consider it during the physical fabrication parameter optimization stage (§ 7).

5.3 Near-Field Major Flow Design

Unlike minor flows, which can be optimized through gradient-based techniques, major flow design requires strategic placement decisions informed by network topology analysis.

Identifying coverage bottlenecks. We begin by analyzing the baseline coverage, *i.e.*, areas reachable by the AP's direct beams without metasurface assistance. Critical bottlenecks are regions that: *(i)* are severely blocked from LoS, thus creating coverage holes; *(ii)* require paths through multiple turns or obstacles; or *(iii)* are geometrically positioned such that no single metasurface can bridge the gap between AP and target region.

For each bottleneck, we compute the link budget deficit, *i.e.*, the gap between required SNR for reliable communication and the best achievable SNR using direct paths or single-hop metasurface relays. *Bottlenecks with large deficits are prioritized for multi-hop major flow construction.* In practice, 1–3 hops are typical since cumulative path loss and efficiency degradation quickly outweigh the focusing gain at longer chains.

Beamforming-aware relay path planning. We propose a passive relay path planning method to identify the candidate metasurface locations that form a major flow. Note that each major flow should originate from the AP, and is limited by the AP's steerable beams (*e.g.*, 8 sectors). For each such beam, we estimate the furthest point reachable with sufficient signal strength, accounting for beamforming gain and free-space path loss (FSPL). From these reachable points, we compute relay paths to the bottleneck regions using a modified shortest-path algorithm, in which the candidate node locations are constrained by available mounting locations of the metasurfaces, and the connectivity between nodes are constrained by the angular FoV of each metasurface.

The link metric is approximated by the FSPL between locations, additionally scaled by *(i)* the far-field reflection beamforming gain, or near-field focusing gain (when metasurfaces are placed at moderate distances); *(ii)* metasurface efficiency factor η. This reduces to finding paths that minimize cumulative distance, transforming multiplicative relay loss into additive distance loss, as discussed in § 4.1.

Near-field focusing optimization. Once major flow paths are identified, we optimize the near-field beam pattern of each metasurface in the relay chain.

For a metasurface positioned at distance d_1 from the AP and distance d_2 from the next relay or interface point, we apply near-field spherical wavefront compensation to focus the reflected beam, where the focal distance is calculated with the focal length equation. The key insight is that the incident wave from a transmitter arrives as a spherical wavefront (not a plane wave) within near-field. To focus the reflection at the next relay point, we compensate for both the spherical incident wavefront and create a converging spherical reflected wavefront. Specifically, the phase at element position r is adjusted by [17]: $\Delta\Phi(r) = k_0(\|r-r_{\text{src}}\|+\|r-r_{\text{next}}\|-(d_1+d_2))$, where r_{src} is the AP position, r_{next} is the next relay point, and $k_0 = 2\pi/\lambda$ is the wavenumber. This compensation transforms the incident spherical wave into a focused spherical wave converging at the next relay point, realizing the "concentration gain" that combats spreading loss.

We optimize this focal distance through gradient descent, which maximizes the rate towards downstream metasurfaces or user locations. Empirically, we find that focusing directly at the next metasurface location yields near-optimal performance for most indoor scenarios. This is because: *(i)* it maximizes power delivery to the relay point, and *(ii)* the next metasurface can then re-focus for its subsequent hop. Fig. 10 validates this finding, showing reduced gain from intermediate focal points.

Interface metasurface selection. The final step of major flow design selects which metasurfaces will serve as interfaces between major and minor flows. These interface metasurfaces must: *(i)* receive strong signals from the major flow relay chain, and *(ii)* be well-positioned to fan out coverage to many receiver locations.

We formulate this as a *maximum coverage problem.* Given the signal strength estimation at each candidate interface location (computed from the major flow paths), we need to select a subset that maximizes the total area coverable with sufficient SNR. We solve this subproblem using a greedy algorithm, which iteratively selects the interface metasurface with highest marginal coverage benefit. *(i)* For each candidate interface location i, estimate the received signal strength RSS_i from the major flow chain (using the optimized near-field focused beams). *(ii)* For each interface candidate i, compute its coverage set C_i, *i.e.*, all receiver locations that can achieve SNR > receiver sensitivity threshold when served by a fan beam from i with power RSS_i. *(iii)* Greedily select interface metasurfaces which achieves a $(1 - 1/e) \approx 0.63$ lower bound to the maximum coverage solution [46].

5.4 Minor Flow Design and Optimization

Unlike major flows, which are sparse and highly directional, minor flows must balance beam spread (for angular diversity) with sufficient gain (for adequate SNR).

Pre-pruning and initialization. Before optimization, we apply physics-informed pruning to eliminate unsuitable minor flow metasurface locations that *(i)* cannot receive strong signals from any interface metasurface or Tx; *(ii)* are geometrically unfavorable (*e.g.*, facing away from uncovered receiver areas); *(iii)* create excessive interference to existing major flows. Note that additional implicit pruning occurs during the actual gradient-based optimization process, since metasurfaces that provide negligible coverage benefit naturally receive selection weights near zero.

In addition, we perform initialization based on each candidate's marginal contribution to coverage, given the already-optimized major flows, which approximates the metasurface's value in completing coverage, and is used to initialize both selection probabilities (higher-scoring locations start with higher selection weights) and beam patterns (initially pointed toward high-value uncovered regions). This warm start helps improve convergence speed and solution quality.

Annealed passive routing optimization (APRO). The core challenge in minor flow optimization is coordinating metasurfaces to cover different regions, avoiding redundant coverage while maintaining sufficient angular robustness. We address this through APRO (Fig. 4), which follows a 4-step iterative optimization process. For each iteration t, APRO *(i)* computes adaptive routing weights $W^{(t)}[m, r]$; *(ii)* evaluates objective using annealed path combination; *(iii)* updates metasurface patterns; *(iv)* updates marginal contribution scores and selection weights α_m. For each component:

(i) Adaptive passive routing. We maintain a routing weight matrix $W[m, r]$ specifying how much metasurface m should prioritize covering user location r. These weights are adapted during optimization to encourage specialization: if metasurface m provides

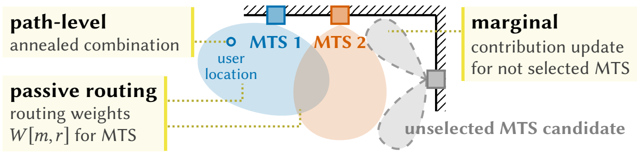

Figure 4: Illustration of annealed passive routing optimization (APRO) for coordinated specialization.

strong coverage to region A while metasurface m' covers region B, their routing weights are adjusted to reinforce this specialization. The routing weights are computed via a soft assignment based on each source's signal contribution:

$$W[m,r] = x_{r,m}^{1/\beta} \Big/ \sum_j x_{r,j}^{1/\beta} \,, \tag{5}$$

where $x_{r,m}$ is the signal contribution from metasurface m to user r (or the direct path for $m = 0$), and $\beta \in (0, 1]$ is a temperature parameter [58]. Initially (small β), the path leans toward a specific metasurface for coverage. As β anneals toward 1, routing softens to true proportional contribution, letting metasurfaces collaborate for multi-path coverage.

(ii) Path-level annealed combination. When computing SNR at each receiver, we use annealed weighting across multipaths rather than simple power summation. Paths with higher gains receive disproportionately higher weights:

$$\mathrm{SNR}_{\mathrm{eff}}(r) = \Big(\sum_m \alpha_m W[m,r] \cdot \mathrm{SNR}_m(r)^{1/\beta} \Big)^{\beta}, \tag{6}$$

where $\mathrm{SNR}_m(r)$ is the contribution from metasurface m. This is governed by the same temperature β as the routing weights. As β anneals toward 1, the combination transitions from winner-takes-all to linear summation, allowing constructive multi-path contributions. This aligns with the *path diversity* rule: each location should have a few strong arrival paths rather than many weak ones.

(iii) Pattern update. We update each metasurface's beam pattern parameters (§ 5.2) via routing-weighted gradient descent through the differentiable channel model. As Proposition 4.2 notes, the relay–coverage separation is not always strict in practice, and therefore major flow and interface metasurface patterns are also refined.

(iv) Marginal contribution updates. After each iteration, we update the metasurface selection parameters by comparing each metasurface's marginal contribution to the current set. Specifically, for each selected metasurface, we compute the coverage loss if it were removed; and vice versa for unselected candidates. Since the routing-weighted objective decouples across metasurfaces, and each α_m is binary, selection under a budget of K reduces to ranking candidates by marginal contribution and selecting the top K. A diversity discount is applied for candidates whose coverage overlaps with already-selected metasurfaces. A broader pool of unselected candidates receives lightweight pattern updates to keep their marginal contribution scores fresh for subsequent iterations.

After convergence of the minor flow metasurface optimization, we run a final pattern refinement stage by freezing the metasurface locations. First, swap refinement evaluates a shortlist of candidate replacements, with dedicated pattern refinement for each candidate configuration. Then, we optimize only the phase patterns of the final selected ones.

Phase finetuning and fabrication optimization. The final stage maps abstract beam parameters to physical phase patterns while accounting for fabrication constraints:

(i) Phase finetuning: With placement and beam patterns fixed, we use gradient descent to optimize the actual phase values of the elements on all the metasurfaces. This allows us to account for coherent combination across coordinated metasurfaces inherently with our wideband site-specific channel confluence model.

(ii) Physical parameter optimization: We map continuous phases to discrete fabrication parameters (*e.g.*, dielectric thickness for 3D-printed metasurfaces) accounting for: *(1)* phase quantization error from limited resolution; *(2)* angle- and frequency-dependent unit cell responses; *(3)* grating lobes and other non-ideal effects.

This physical optimization uses late acceptance hill-climbing (LAHC) combined with local search, as gradient information is no longer reliable due to discrete parameters. Starting from the quantized version of the optimized continuous phases, LAHC explores local modifications that improve robustness to fabrication variations and wideband performance.

Improving scalability utilizing low-rank manifold. The hierarchical major–minor decomposition is critical for scalability: major flow optimization runs once to establish the backbone, then minor flow optimization needs only to complete coverage rather than solve the full end-to-end problem. This reduces complexity from exponential to approximately linear in the number of minor flow metasurfaces.

For very large environments, we additionally employ an adaptive low-rank manifold grouping method that exploits the inherent low-rank structure. Our key insight is that candidate metasurface patterns exhibit strong spatial correlations—nearby metasurfaces (similarity adaptively calculated based on geometric location and path information) share similar propagation characteristics and work toward common coverage objectives, suggesting the true degrees of freedom are far fewer than the raw parameter count. Thus, within APRO, we can update the unselected metasurfaces in batch, as they are correlated for marginal contribution.

6 Handling Network Dynamics and Interference

The major–minor flow topology and hierarchical optimization framework produce a static passive metasurface network with fixed beam pattern configurations. A key question remains: *how does this static infrastructure adapt to dynamic environments with mobile users, time-varying channels, and multi-user interference?* We now elaborate on the solutions.

6.1 Seamless Integration with Standard Network Stack

FlowForm passively modifies the RF environment, operating transparently to standard mmWave radios such as IEEE 802.11ad [24] and 802.11ay [25]. To these radios, FlowForm metasurfaces function as natural reflectors that redirect and reshape incident beams. Standard mmWave MAC protocols handle channel adaptation through beam scanning. For example, 802.11ad AP performs periodic sector-level sweep (SLS), transmitting training frames across predefined beam directions (typically 8–32 beams). Mobile devices respond with their best received beam, and the AP selects the beam pair that maximizes SNR [24]. FlowForm enhances this standard process: the metasurface network creates additional paths between the AP and users, which manifest as stronger RSS along specific beam directions. The AP discovers these metasurface-enhanced paths implicitly through its beam scanning. No explicit knowledge of the metasurfaces is required. When a mobile user moves, standard beam scanning automatically selects new AP beam directions that may leverage different metasurfaces or different paths.

6.2 Robustness to Channel Dynamics

Diversity scoring during major–minor flow optimization. Standard mmWave radios have limited field-of-view (FoV), typically 120° [64], due to the directional nature of their antenna or phased array panel. If coverage relies on a single strong path from one direction, a modest device rotation can completely block the signal. To ensure robustness, we incorporate diversity scoring into the optimization objective in § 5, which encourages multiple beams for the same user location. Specifically, we compute the AP and user diversity score for each user location:

$$\begin{cases} S_{d,r} = N_s^{1-\frac{1}{\alpha_r}} \cdot \left(\sum_{i=1}^{N_s} \bar{r}_i^{\alpha_r}\right)^{\frac{1}{\alpha_r}} \in (0,1], \\ S_{d,t} = N_b^{1-\frac{1}{\alpha_t}} \cdot \left(\sum_{b=1}^{N_b} \bar{r}_b^{\alpha_t}\right)^{\frac{1}{\alpha_t}} \in (0,1], \end{cases} \tag{7}$$

where $\bar{r}_i$ is the normalized rate of sector i (divided by the sum of all N_s sector rates, with zero clamped cosine projection representing the antenna element pattern), $\alpha_{\{t,r\}} \in (0,1]$ is the diversity exponent, with lower values encouraging more diversity, and $\bar{r}_b$ is the normalized rate of beam b from the AP (divided by the sum of all N_b beam rates), after applying the user diversity score. As $\alpha_{\{t,r\}}$ decreases, the diversity score increasingly favors scenarios where multiple sectors provide comparable rates rather than one dominant sector. In practice, performance is insensitive to moderate variations of α, and we find $\alpha_{\{t,r\}} = 0.8$ to be a good tradeoff between diversity and average rate. Finally, the diversity scores are incorporated into the optimization objective (§ 5.1) by weighing the estimated rate at each location by $S_{d,r}$. The diversity scoring mechanism ensures multiple metasurface relay paths serve users at diverse locations and orientations. As users move or rotate, the AP's periodic beam scanning discovers and switches to the best beam via a new metasurface relay path.

Long-term environmental changes. Since mmWave propagation is dominated by LoS and a small number of strong reflectors [62], minor furniture rearrangement rarely disrupts the dominant paths. FlowForm further reduces sensitivity to such changes by relying on engineered metasurface paths rather than ambient background reflections. For more substantial environmental changes, such as semi-permanent obstacles, the metasurface network remains effective because the optimization (§ 5) explicitly accounts for robustness. First, during the design-time optimization (§ 5.4), paths are weighted by credibility scores based on geometric stability and likelihood of dynamic blockage. So the metasurface placement naturally favors more stable propagation paths. Second, the minor flow design creates overlapping coverage from multiple metasurfaces, which adds to the path diversity. Third, if major environmental changes occur (*e.g.*, office reconfiguration), FlowForm optimization can be re-run using an updated environment model.

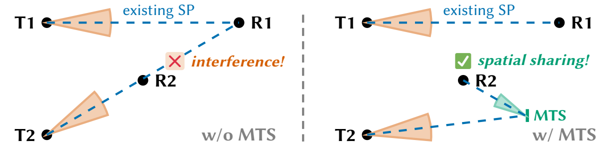

Figure 5: Illustration of metasurface-aided spatial sharing by providing alternative paths.

6.3 Multi-User Interference and Spatial Sharing

Multi-AP coordination. FlowForm metasurface network is fully compatible with standard time-division multiple access (TDMA)-based interference coordination protocols. Unlike active relays or RIS, which must be explicitly assigned to serve specific APs, passive metasurfaces simply reflect incident energy according to their designed patterns, naturally enhancing active AP's transmission without requiring synchronization. For users, mobile devices automatically associate with the AP providing the highest RSS as in 802.11ad protocol [24], and FlowForm's RSS enhancement requires no protocol modifications.

Enhancing spatial sharing. Prior measurement studies demonstrated that even directional beams create substantial interference through sidelobes [48]. With metasurfaces, interference can occur even for spatially-separated Tx–Rx pairs if a metasurface reflects signals from one pair toward the other pair's Rx. The standard MAC protocols define spatial sharing mechanisms to enable concurrent transmissions [24, 25]. An AP can establish a service period (SP) for one user while another SP is already active, provided two conditions are met: *(i)* the candidate SP is strong enough compared to noise and interference from the existing SP; *(ii)* the existing SP receiver is not interfered by the candidate SP.

FlowForm can enhance spatial sharing by providing more beam options (*i.e.*, different paths through different sequences of metasurfaces), which increases the likelihood that spatially-separated beam pairs can be found, as illustrated in Fig. 5. To maximize such opportunities, we augment the major–minor flow optimization to become interference-aware. Consider the downlink case for simplicity (uplink analysis follows similarly). Let $\mathrm{RSS}_{i(k)\to j}$ denote RSS from AP i with beam k to user j, and $I_{i(k)\to j}$ the aggregate interference signal strength at user j. Concurrent SPs can be established when:

$$\mathrm{SINR}_{i(k)\to j} = \mathrm{RSS}_{i(k)\to j} \big/ (N + I_{i(k)\to j}) > \gamma, \tag{8}$$

where N is the noise power and γ is an implementation-dependent threshold.

We update the optimization objective (§ 5.1) by using the SINR in Eq. (8) instead of the original SNR, across all possible user locations. The resulting metasurface location and beam configuration will thus take into account the spatial sharing among the multiple APs (assuming known locations). Joint optimization of AP placement and metasurface configuration is a natural extension that plugs directly into this interference-aware objective.

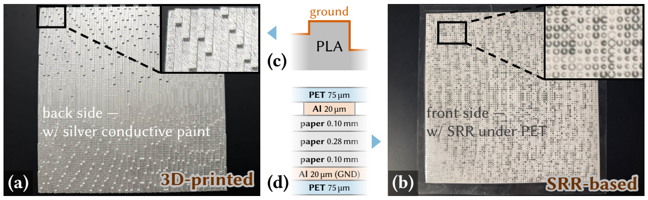

Figure 6: Low-cost and scalable metasurface fabrication in FlowForm. (a, c) 3D-printed metasurface [53] with PLA substrate and conductive paint ground. (b, d) SRR-based metasurface [41] with hot-stamped aluminum (Al) foil on paper substrate.

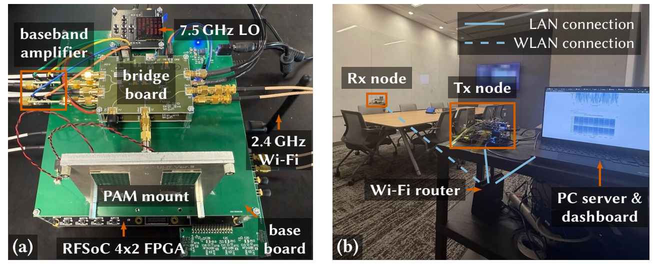

Figure 7: Experimental setup: (a) mmWave software radio for measurements and validation. (b) Example deployment in a conference room.

7 Implementation

Site-specific channel confluence model. FlowForm's optimization (§ 5) requires repeatedly evaluating how different metasurface placements and configurations $\boldsymbol{\Psi}$ affect coverage and throughput. To meet this requirement, we develop a differentiable ray tracing model [11] to compute geometric propagation paths between all nodes (Tx, Rx, and MTS). Each path has a complex path gain, which is scaled by path loss, Tx/Rx beamforming gain and metasurface reflection patterns. Our ray tracer extends the open-source Sionna 1.2.1 [4, 20], by modeling the frequency-dependent signal propagation and multipath across the 2.16 GHz 802.11ad/ay bandwidth.

Ray tracing requires a 3D geometric model of the deployment environment annotated with electromagnetic material properties. Since FlowForm's optimization framework implicitly constrains the dominant signal energy within the metasurface network, it is insensitive to the ambient environmental structures (§ 8.6). Therefore, we build a simplified workflow to create the environmental model as an input to the ray tracing model. Specifically, for each deployment site, we leverage readily-available 2D floor plans with semantic labels (*e.g.*, walls, doors, windows) to synthesize a 3D environment, augmented with height information and furniture to minimize the sim-to-real gap. We then manually annotate the dominant objects' electromagnetic properties (based on [26]) to approximately match the actual environment.

Passive metasurface fabrication. FlowForm metasurfaces are fabricated using low-cost 3D printing that enables rapid prototyping and deployment at scale, as shown in Fig. 6. Each metasurface consists of a dielectric substrate with conductive ground plane following [53]. Each unit cell is a half-wavelength dielectric cuboid with height h that determines the phase shift. The phase-shift is in turn determined by the major-minor flow optimization (§ 5). In practice, we find that the quality of the fabricated metasurface beam pattern is impacted by several fabrication artifacts, including *(i)* phase quantization error due to limited fabrication resolution; *(ii)* amplitude response variation of unit cells; *(iii)* angle and frequency-dependent response of unit cells. These non-idealities are partially accounted for in optimization through the efficiency parameter η (§ 5.2). To achieve an even more accurate phase-to-h mapping, we apply a two-stage physical finetuning.

First, for major flow metasurfaces designed with near-field focusing (§ 5.3), the beam-level optimization assumes a specific focal distance d_{foc}. However, the distribution of target user areas can change after joint optimization and the unit-cell non-idealities including discrete height quantization may shift the effective focal point. We thus re-optimize d_{foc} via gradient descent, which searches for the quantized heights h that maximizes average rate across user locations (Alg. 2 in Appendix C.3). We find this can typically recover about half of the loss from fabrication artifacts.

Second, we employ late acceptance hill-climbing (LAHC) [9] to refine the individual element heights. LAHC explores the discrete solution space through local modifications, accepting improvements based on a history of recent objective values (Alg. 3 in Appendix C.4). This optimization further improves wideband performance, with the combined physical optimization (PO) advantage shown in Fig. 11(a).

mmWave network testbed. FlowForm is designed to work with standard, off-the-shelf 802.11ad/ay devices without modification, requiring no per-beam or channel state feedback from the AP at runtime. However, commercial 802.11ad/ay devices do not expose detailed channel information (*e.g.*, per-beam RSS and wideband channel impulse response (CIR)) needed for fine-grained validation of the metasurface network design. We thus build a mmWave software defined radio (SDR) testbed (Fig. 7(a)) for measurement and validation purposes. Our testbed extends the M-Cube SDR [67, 69] as follows: ***(i)*** We implement a real-time CIR estimation module, using the RFSoC 4x2 FPGA board [57] as baseband processor which provides 5 GSps sampling rate. ***(ii)*** We repurpose the 8×8 phased arrays from 802.11ay standard-compatible Mikrotik Cube 60Pro routers [44], and use the GPIO on RFSoC for real-time beam configuration. ***(iii)*** We implement a remote control server that can configure the Tx/Rx radios' beam patterns in real-time and collect per-beam CIR.

8 Evaluation

8.1 Experimental Setup

We deploy FlowForm across 5 indoor environments with varying complexity (Fig. 19). Scene A is a small hallway (69 m^2) with minimal obstacles. Scenes B–C are medium-sized office floors (133–325 m^2) with furniture and partitions. Scene D and E are large corridor areas (383 m^2 and 462 m^2) with severe blockage from columns and walls.

Baselines and metasurface configuration. For system-level evaluation, we compare FlowForm against 3 baselines:

(i) No MTS, which employs only direct AP-to-user paths with standard 802.11ad beamforming;

(ii) AutoMS+, an extended version of AutoMS [41] (§ 2), where we apply the same wideband modeling and receiver FoV constraints

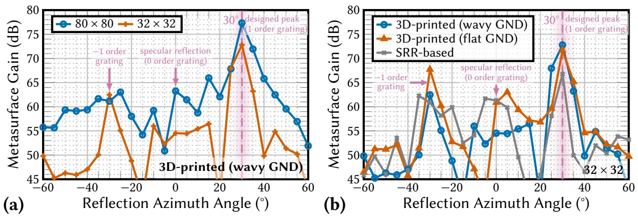

Figure 8: Measured beam patterns w/ 30° designed steering; y-axis is gain *w.r.t.* a single unit-cell element. (a) Different metasurface sizes. (b) 3D-printed vs. SRR-based metasurfaces.

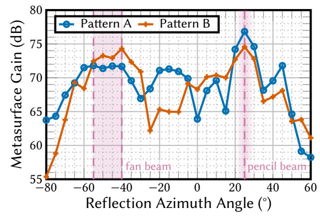

Figure 9: Measured metasurface beam pattern with 2-arm quadratic patterns.

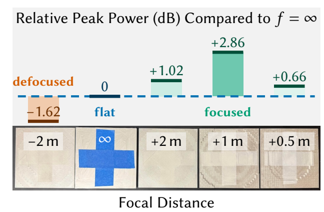

Figure 10: Near-field focusing peak power with different focal distances.

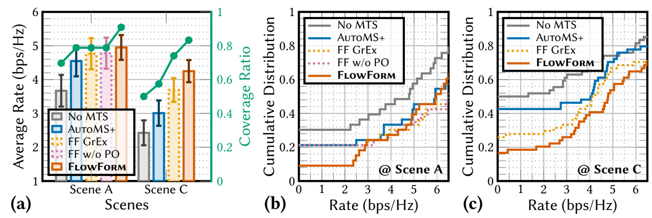

Figure 11: Measurement comparisons with different optimization approaches. (a) Average rate comparison. (b, c) Cumulative distribution of rate.

as in FlowForm (§ 5.4), so as to achieve a fair comparison excluding the hardware artifacts.

(iii) FF GrEx, which exhaustively searches metasurface placements and greedily adds metasurfaces with best marginal coverage. This represents a strong baseline on what can be achieved without the major–minor hierarchy, though it does not scale to large deployments.

Unless otherwise specified, we use 80 × 80 element 3D-printed metasurfaces with wavy ground plane design [53]. Each metasurface measures 20 cm × 20 cm with 6400 elements and costs approximately $2 to fabricate (Appendix C.2).

Measurement methodology. For each test scenario, we measure: ***(i)*** Per-beam CIR, more specifically the 256-tap CIR for all 8 transmit beams at each user location, which provides detailed multipath information. RSS is then extracted from CIR peak power. ***(ii)*** Physical rate, calculated using the Shannon capacity formula [60] considering the radio hardware implementation loss. ***(iii)*** Coverage ratio, defined as the percentage of test locations achieving RSS > -68 dBm (receiver sensitivity for 802.11ad/ay MCS 1, [24, 25]). We measure 10–25 random uniformly distributed user locations in each environment, with 3 user orientations per location (0°, 120°, 240°).

8.2 Micro Benchmarks

We first validate that individual metasurfaces can achieve their designed beam patterns with minimal hardware artifacts.

Single-beam steering. Fig. 8(a) shows measured beam patterns with different metasurface sizes. The 80×80 metasurface achieves ~5 dB higher peak gain than 32 × 32, confirming the size-gain relationship. Both metasurfaces achieve the designed steering angle (30°) within 3° error. Fig. 8(b) reveals a non-ideal beamforming effect. The main lobe (*i.e.*, designed +1 order) contains ~70% of the reflected power. This efficiency $\eta \approx 0.7$ matches the parameter used in our metasurface abstraction model (§ 5.2). In addition, the 3D-printed metasurfaces with a wavy ground plane [53] achieves the highest peak gain with smaller specular reflection (0 order grating) and conjugate lobe (−1 order grating), in comparison to flat ground 3D-printed metasurfaces [65] and SRR-based metasurfaces [41]. We thus use the wavy ground design by default in our experiments.

Metasurface beam pattern abstraction. Fig. 9 demonstrates the multi-arm quadratic beam representation (§ 5.2) working in practice. We fabricate a 80 × 80 metasurface with 2-arm pattern: a narrow pencil beam for major flow relay and a wider fan beam for minor flow coverage. Beam Pattern A emphasizes the fan beam (70% power), while Pattern B emphasizes the pencil beam (70% power). Measured patterns match the design with the relative power distribution preserved. This validates compact parametric representation captures essential beam characteristics while maintaining physical practicality.

Near-field focusing validation. To verify the near-field focusing gain (§ 5.3), we place 32 × 32 metasurfaces 2 m from the Tx and Rx (within near-field region) with various focal distances. When the metasurface is focused at the designed distance (1 m from the metasurface), we observe 2.86 dB RSS gain compared to a flat (non-focused) pattern. Over-focusing ($f = +0.5$ m) or under-focusing ($f = +2$ m) both degrade performance, confirming that the focal distance optimization (§ 5.3) is critical.

8.3 Coverage Expansion

Coverage expansion performance. Fig. 11(a) shows the measured average rate across user locations and coverage ratio in two scenes, where error bars indicate standard error across measurements. With 16 metasurfaces (5 major + 11 minor), FlowForm achieves 75% rate improvement and 67% coverage improvement over no MTS in scene C. Despite using the same number of metasurfaces, FlowForm outperforms AutoMS+ by 9% even with minor flow only optimization in a small environment (scene A). In scene C, FlowForm leverages metasurface coordination to create concentrated, specialized patterns, outperforming AutoMS+ by 41% and FF GrEx by 15%, demonstrating the major-minor hierarchy is near-optimal yet far more scalable. The rate gains consistently exceed the standard error across all scenes, and the CDFs in Fig. 11(b, c) show that FlowForm particularly benefits originally poorly covered locations, lifting the lower tail of the rate distribution, as further illustrated by the per-location breakdown in Fig. 12.

Scaling with number of metasurfaces. Fig. 13(a) shows performance vs. number of deployed metasurfaces in scene C. With 6 metasurfaces (2 major + 4 minor), average rate improves by 47%

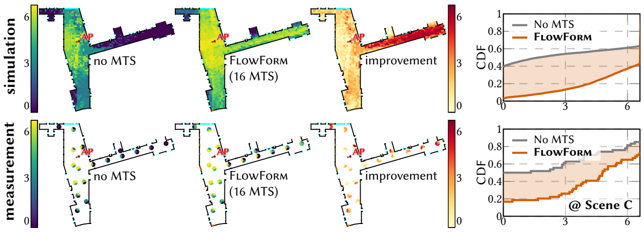

Figure 12: Measured and simulated coverage expansion showing consistent rate (bps/Hz) results.

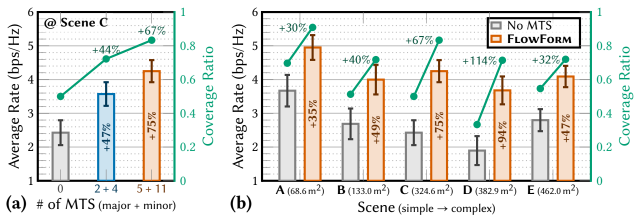

Figure 13: Measurement results comparisons of average rate and coverage ratio. (a) Different number of metasurfaces. (b) Different scenes.

and coverage by 44%. Increasing to 16 metasurfaces (5 major + 11 minor) provides further gains with diminishing returns: initial metasurfaces address critical bottlenecks while more additions fill remaining gaps.

Performance across environments. As shown in Fig. 13(b), FlowForm's performance gain initially grows with scene complexity (average rate improved by 35% in A to 75% in C) due to larger margin for enhancement. However, in very large or geometrically unfavorable scenes (E), FlowForm gains are constrained by propagation loss. Scene D, despite low baseline coverage, benefits significantly from multi-hop relaying that effectively shortens propagation distance, showing 94% rate and 114% coverage improvement. In general, FlowForm's gain is largest when NLoS regions dominate and sufficient wall surfaces are available for metasurface mounting (*e.g.*, scene C and D). Gains are smaller when LoS already provides adequate coverage (*e.g.*, scenes B and E) or metasurface deployment is infeasible to improve NLoS due to site limitations.

FoV and orientation robustness. Extending the results in Fig. 11, we show the measured per-direction rate in Fig. 12. Without FlowForm, 120° rotation causes rate to drop by 3 bps/Hz (often causing complete outage). With FlowForm, for those locations the same rotation causes only 1–2 bps/Hz variation because diverse arrival angles (from different metasurfaces) ensure at least one strong path remains within the receiver's FoV. This directly validates the diversity mechanism in § 6.2.

8.4 Robustness to Dynamics and Interference

Receiver mobility and orientation dynamics. Fig. 14 shows RSS for 8 AP beams as a user moves through Scene C. Without metasurfaces, RSS varies ~20 dB with simultaneous dropouts. FlowForm reduces variation to ~5 dB while maintaining multiple beams above −50 dBm, so that automatic beam selection can maintain stable connectivity. Fig. 14(b) shows a more challenging scenario: a user

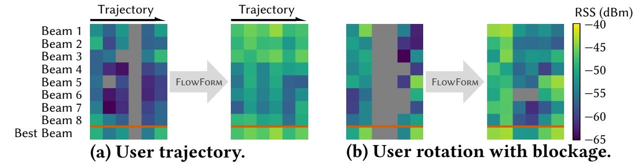

Figure 14: Channel dynamics w/o and w/ FlowForm. (a) User trajectory. (b) User rotation with blockage.

rotates with blocked LoS, emulating typical mobile device use. Without metasurfaces, the user frequently experiences complete outage for rotation, whereas FlowForm maintains a more stable link. In 802.11ad, frequent beam scanning (every 100 ms) is necessary to track the best beam as channels vary [14]. With FlowForm's stable, diverse coverage, we find that scanning intervals can be extended to 250–300 ms without performance degradation, which recovers 60% of the scanning overhead for data transmission. This validates the claim that robust coverage reduces beam management overhead (§ 6.2). We further interpolate the same trajectory at finer granularity in simulation to extract link-level metrics: FlowForm achieves 6.36 bps/Hz mean rate with no blackout and 0.14 bps/Hz jitter, vs. 3.88 bps/Hz, 18% blackout, and 0.54 bps/Hz jitter without metasurfaces.

Beyond short-term mobility, we observe across our coverage experiments that normal furniture movement (*e.g.*, chairs) has negligible impact on measured RSS, consistent with the sparsity of mmWave channels dominated by LoS and major reflectors. FlowForm's reliance on engineered metasurface paths further limits sensitivity to such environmental changes (§ 6.2), and redeployment is only necessary when multiple major flow paths are permanently obstructed.

Multi-user interference and spatial sharing. Fig. 15 evaluates spatial sharing performance with two concurrent AP-receiver pairs in scene C. We compare against the average throughput (mapped from RSS [25]) before and after FlowForm deployment over three scheduling policies: *(P1)* no spatial sharing (TDMA), *(P2)* spatial sharing with one AP priority (maximizing one AP's throughput), and *(P3)* maximum sum throughput. The triangle points for P1 and P2 denote maximum possible rate without fairness.

Extending coverage with metasurfaces could potentially introduce additional interference, degrading spatial sharing performance. However, Fig. 15 shows that FlowForm's SINR-aware optimization (§ 6.3) avoids this: the spatial sharing gains (P2 and P3) remain consistent before and after FlowForm deployment (~43% in both cases compared to no MTS) confirming that coverage extension does not come at the cost of spatial reuse efficiency. As the number of users scales, time sharing follows standard 802.11ad/ay MAC scheduling. FlowForm's improved spatial sharing enables more concurrent transmissions, so per-user throughput degrades more gracefully than without metasurfaces.

8.5 Comparison With Active Metasurfaces

Since real mmWave RIS panels are cost prohibitive, we simulate RIS performance using our ray tracing framework (§ 7). For each user location and orientation, we optimize the RIS' element-wise phase reconfiguration via gradient descent to maximize SNR. This

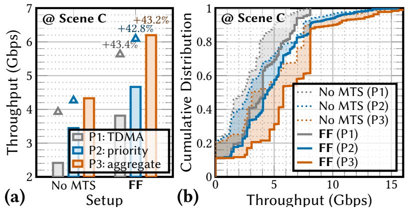

Figure 15: Throughput gain from FlowForm with spatial sharing. (a) Comparison. (b) Distribution.

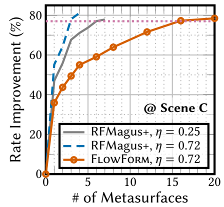

Figure 16: Scaling comparison with RFMagus+.

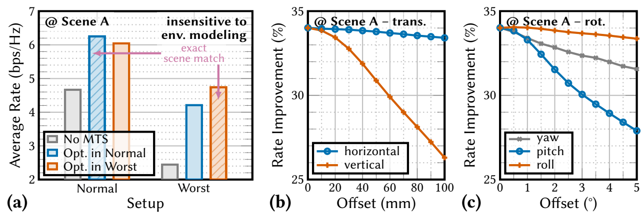

Figure 17: Sensitivity analysis of FlowForm in simulation. (a) Sensitivity to environmental modeling. (b, c) Sensitivity to metasurface placement.

represents an optimistic upper bound on RIS performance, as practical RIS would additionally suffer from hardware imperfections and control overhead. We call this baseline *RFMagus+* because it extends RFMagus [34] with joint optimization across multiple RIS surfaces rather than alternating optimization. We compare RIS with two hardware efficiency levels: $\eta = 0.72$ (matching our passive metasurfaces) and $\eta = 0.25$ (realistic for 60 GHz RIS [47, 59], which suffer from high insertion loss).

Fig. 16 shows the results in scene C. Compared with the RIS with a realistic hardware efficiency ($\eta = 0.25$), FlowForm needs only 2.7× more metasurface units (16 vs. 6) for equivalent coverage. Yet, the cost of the RIS deployment at 60 GHz band amounts to $6k+ minimum, in contrast to only $32 for FlowForm (~200× cost reduction even excluding the cost of deployment, power, control channel, *etc.*).

8.6 Sensitivity Analysis

Sensitivity to environmental modeling. We proceed to evaluate robustness to modeling errors and deployment imperfections. Fig. 17(a) compares performance when the test environment differs from the optimization environment. We create two scenarios: *(i)* Normal: Environment model matches the actual deployment. *(ii)* Worst case: All objects modeled as perfect absorbers (no environmental reflection), representing severe modeling errors.

FlowForm performance degrades only ∼11% in the worst-case scenario. When optimized in the "normal" environment but deployed in "worst," the average rates are still more than 70% better than no MTS. When optimized in the "worst" environment but deployed in "normal," it performs nearly as well. This robustness stems from FlowForm's design philosophy (§ 4): the major–minor flow topology prioritizes strong, stable LoS and single-reflection paths over weak multi-bounce paths that heavily depend on environmental modeling. Because the dominant LoS propagation mechanisms are captured even in simplified models, errors in environment modeling have limited impact on performance.

Sensitivity to metasurface placement. Fig. 17(b, c) evaluates sensitivity to metasurface placement errors. We simulate deploying metasurfaces with random offsets from their optimized positions and measure resulting rate degradation. FlowForm tolerates ±10 cm horizontal offset with < 1% rate degradation. Vertical offset is more sensitive due to narrow beamwidth: ±3 cm causes ∼1% degradation, while ±10 cm causes ∼8% loss. Yaw (horizontal rotation) is tolerant due to wider beamwidth, while pitch (vertical tilt) is most sensitive (±3° causes ∼10% degradation). Roll has minimal impact due to metasurface symmetry. Since typical wall construction tolerances are within 0.5° [3, 15], mounting errors within this range would not exhibit noticeable performance loss, confirming that manual installation using measured floor plans is practical.

8.7 Complexity Analysis

FlowForm's computational workload occurs only during major–minor flow optimization (§ 5) before deployment. With an Nvidia A6000 GPU, optimization takes 10 minutes for simple scenes (scene A) and up to 1 hour for complex ones (scene E). The metasurface abstraction framework and major–minor decomposition significantly reduce complexity compared to AutoMS [41] (∼8 hours for scene E with worse performance) and FF GrEx (multiplied in time usage).

9 Conclusion

This work introduces FlowForm, a system designed to achieve scalable passive metasurface networking for mmWave communications. Our approach addresses critical challenges in 60 GHz coverage expansion through a sophisticated optimization framework that jointly manages metasurface placement and phase design. By leveraging a reliability-aware differentiable ray tracer, our system explicitly accounts for real-world factors like multi-AP interference and channel dynamics, creating a robust and transparent solution. Experiments across five indoor environments validate substantial gains at orders of magnitude lower cost than active RIS alternatives. Ultimately, FlowForm provides a practical pathway to engineering the radio environment itself, transforming unreliable mmWave links into a dense and resilient communication network.

Ethics statement. FlowForm does not raise any ethical issues. No human subject is involved.

Acknowledgments

We appreciate the insightful comments and feedback from the anonymous Reviewers and Shepherd. We also appreciate the helpful discussions and experimental assistance from Justin Chen. This research was sponsored by the OUSD (R&E)/RT&L and was accomplished under Cooperative Agreement Number W911NF-20-2-0267. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the ONR, ARL and OUSD(R&E)/RT&L or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. This research was also supported in part by NSF under Grant CNS-2312715.

References

[1] 5G Americas and Small Cell Forum. 2019. *Precision Planning for 5G Era Networks with Small Cells.* White Paper SCF 230.10.01. 5G Americas & Small Cell Forum. Retrieved June 26, 2026 from https://www.5gamericas.org/wp-content/uploads/2019/10/SCF-Precision-Planning-WP-FINAL-1.pdf

[2] Omid Abari, Dinesh Bharadia, Austin Duffield, and Dina Katabi. 2017. Enabling High-Quality Untethered Virtual Reality. In *14th USENIX Symposium on Networked Systems Design and Implementation* (Boston, MA, USA) *(NSDI '17)*. USENIX Association, Berkeley, CA, USA, 531–544. https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/abari

[3] ACI Committee 117. 2010. *Specification for Tolerances for Concrete Construction and Materials (ACI 117-10) and Commentary (ACI 117R-10).* ACI Standard 117-10. American Concrete Institute. Retrieved June 26, 2026 from https://www.pa.gov/content/dam/copapwp-pagov/en/dli/documents/ucc/documents/2018-icc-code-review-public-comments/szoke-117-10.pdf

[4] Fayçal Aït Aoudia, Jakob Hoydis, Merlin Nimier-David, Sebastian Cammerer, and Alexander Keller. 2025. Sionna RT: Technical Report. https://arxiv.org/abs/2504.21719. arXiv:2504.21719 [cs.IT]

[5] Bambu Lab. 2026. *PLA Basic.* Retrieved June 26, 2026 from https://us.store.bambulab.com/products/pla-basic-filament

[6] Dinesh Bharadia and Sachin Katti. 2014. FastForward: fast and constructive full duplex relays. In *Proceedings of the 2014 ACM Conference on SIGCOMM* (Chicago, Illinois, USA) *(SIGCOMM '14)*. Association for Computing Machinery, New York, NY, USA, 199–210. doi:10.1145/2619239.2626327

[7] Sudip Biswas, Satyanarayana Vuppala, Jiang Xue, and Tharmalingam Ratnarajah. 2015. On the Performance of Relay Aided Millimeter Wave Networks. *IEEE J. Sel. Topics Signal Process.* 10, 3 (2015), 576–588. doi:10.1109/JSTSP.2015.2509951

[8] Michael Boyarsky, Mohammadreza F Imani, and David R Smith. 2020. Grating lobe suppression in metasurface antenna arrays with a waveguide feed layer. *Opt. Express* 28, 16 (2020), 23991–24004. doi:10.1364/OE.398440

[9] Edmund K Burke and Yuri Bykov. 2017. The late acceptance Hill-Climbing heuristic. *Eur. J. Oper. Res.* 258, 1 (2017), 70–78. doi:10.1016/j.ejor.2016.07.012

[10] Caswell. 2025. *Copper Conductive Paint - 4OZ.* Retrieved June 26, 2026 from https://caswellplating.com/copper-conductive-paint-4oz.html

[11] Xingyu Chen, Xinyu Zhang, Kai Zheng, Xinmin Fang, Tzu-Mao Li, Chris Xiaoxuan Lu, and Zhengxiong Li. 2026. Physically Accurate Differentiable Inverse Rendering for Radio Frequency Digital Twin. arXiv:2603.18026 [eess.SP] https://arxiv.org/abs/2603.18026

[12] Kun Woo Cho, Mohammad H Mazaheri, Jeremy Gummeson, Omid Abari, and Kyle Jamieson. 2023. mmWall: A Steerable, Transflective Metamaterial Surface for NextG mmWave Networks. In *20th USENIX Symposium on Networked Systems Design and Implementation* (Boston, MA, USA) *(NSDI '23)*. USENIX Association, Berkeley, CA, USA, 1647–1665. https://www.usenix.org/conference/nsdi23/presentation/cho-kun-woo

[13] Tie Jun Cui, David R Smith, and Ruopeng Liu. 2010. *Metamaterials: Theory, Design, and Applications.* Springer, New York, NY, USA. doi:10.1007/978-1-4419-0573-4

[14] Mohammed Dahhani, Gentian Jakllari, and André-Luc Beylot. 2019. Association and Reliability in 802.11ad Networks: An Experimental Study. In *2019 IEEE 44th Conference on Local Computer Networks* (Osnabrueck, Germany) *(LCN '19)*. IEEE, New York, NY, USA, 398–405. doi:10.1109/LCN44214.2019.8990684

[15] Danish Standards. 2010. *Execution of concrete structures.* DS/EN Standard 13670:2010. Danish Standards (DS). Retrieved June 26, 2026 from https://webstore.ansi.org/standards/ds/dsen136702010

[16] Manideep Dunna, Chi Zhang, Daniel Sievenpiper, and Dinesh Bharadia. 2020. ScatterMIMO: enabling virtual MIMO with smart surfaces. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking* (London, United Kingdom) *(MobiCom '20)*. Association for Computing Machinery, New York, NY, USA, Article 10, 14 pages. doi:10.1145/3372224.3380887

[17] Chao Feng, Xinyi Li, Yangfan Zhang, Xiaojing Wang, Liqiong Chang, Fuwei Wang, Xinyu Zhang, and Xiaojiang Chen. 2021. RFlens: metasurface-enabled beamforming for IoT communication and sensing. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking* (New Orleans, LA, USA) *(MobiCom '21)*. Association for Computing Machinery, New York, NY, USA, 587–600. doi:10.1145/3447993.3483238

[18] Corentin Fonteneau, Matthieu Crussière, and Bruno Jahan. 2021. A Systematic Beam Broadening Method for Large Phased Arrays. In *2021 Joint European Conference on Networks and Communications & 6G Summit* (Porto, Portugal) *(EuCNC/6G Summit '21)*. IEEE, New York, NY, USA, 7–12. doi:10.1109/EuCNC/6GSummit51104.2021.9482595

[19] Roger A Horn and Charles R Johnson. 2012. *Matrix Analysis* (2 ed.). Cambridge University Press, New York, NY, USA. doi:10.1017/CBO9780511810817

[20] Jakob Hoydis, Sebastian Cammerer, Fayçal Ait Aoudia, Avinash Vem, Nikolaus Binder, Guillermo Marcus, and Alexander Keller. 2023. Sionna: An Open-Source Library for Next-Generation Physical Layer Research. https://arxiv.org/abs/2203.11854. arXiv:2203.11854 [cs.IT]

[21] Kai-Cheng Hsu, Kate Ching-Ju Lin, and Hung-Yu Wei. 2016. Full-duplex delay-and-forward relaying. In *Proceedings of the 17th ACM International Symposium on Mobile Ad Hoc Networking and Computing* (Paderborn, Germany) *(MobiHoc '16)*. Association for Computing Machinery, New York, NY, USA, 221–230. doi:10.1145/2942358.2942370

[22] Chongwen Huang, Alessio Zappone, George C Alexandropoulos, Mérouane Debbah, and Chau Yuen. 2019. Reconfigurable Intelligent Surfaces for Energy Efficiency in Wireless Communication. *IEEE Trans. Wireless Commun.* 18, 8 (2019), 4157–4170. doi:10.1109/TWC.2019.2922609

[23] iBwave. 2026. *iBwave Mobile Survey: Wireless Network Survey Solution.* Retrieved June 26, 2026 from https://www.ibwave.com/ibwave-mobile-survey/

[24] IEEE. 2012. *IEEE 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.* IEEE Standard 802.11ad. IEEE. doi:10.1109/IEEESTD.2012.6392842

[25] IEEE. 2021. *IEEE 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 2: Enhanced Throughput for Operation in License-exempt Bands above 45 GHz.* IEEE Standard 802.11ay. IEEE. doi:10.1109/IEEESTD.2021.9502046

[26] International Telecommunication Union. 2015. *P.2040 : Effects of building materials and structures on radiowave propagation above about 100 MHz.* ITU Recommendation P.2040-1. ITU-R. Retrieved June 26, 2026 from https://www.itu.int/rec/R-REC-P.2040

[27] Ish Kumar Jain, Raghav Subbaraman, and Dinesh Bharadia. 2021. Two beams are better than one: Towards Reliable and High Throughput mmWave Links. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference* (Virtual Event, USA) *(SIGCOMM '21)*. Association for Computing Machinery, New York, NY, USA, 488–502. doi:10.1145/3452296.3472924

[28] Keysight. 2026. *How to Perform 5G Private Network Site Survey.* Retrieved June 26, 2026 from https://www.keysight.com/us/en/use-cases/perform-5g-private-network-site-survey.html

[29] Wahab Khawaja, Ozgur Ozdemir, Yavuz Yapici, Fatih Erden, and Ismail Guvenc. 2020. Coverage Enhancement for NLOS mmWave Links Using Passive Reflectors. *IEEE Open J. Commun. Soc.* 1 (2020), 263–281. doi:10.1109/OJCOMS.2020.2969751

[30] Jacek Kibiłda, Allen B MacKenzie, Mohammad J Abdel-Rahman, Seong Ki Yoo, Lorenzo Galati Giordano, Simon L Cotton, Nicola Marchetti, Walid Saad, William G Scanlon, Adrian Garcia-Rodriguez, David López-Pérez, Holger Claussen, and Luiz A. DaSilva. 2020. Indoor Millimeter-Wave Systems: Design and Performance Evaluation. *Proc. IEEE* 108, 6 (2020), 923–944. doi:10.1109/JPROC.2020.2989189

[31] Minseok Kim, Namjo Ahn, and Song Min Kim. 2024. NR-surface: NextG-ready μW-reconfigurable mmWave metasurface. In *Proceedings of the 21st USENIX Symposium on Networked Systems Design and Implementation* (Santa Clara, CA, USA) *(NSDI '24)*. USENIX Association, Berkeley, CA, USA, 1641–1657. https://www.usenix.org/conference/nsdi24/presentation/kim

[32] Harold W Kuhn and Albert W Tucker. 1951. Nonlinear Programming. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 2. University of California Press, Berkeley, CA, USA, 481–492.

[33] Xinyi Li, Chao Feng, Xiaojing Wang, Yangfan Zhang, Yaxiong Xie, and Xiaojiang Chen. 2023. RF-Bouncer: A Programmable Dual-band Metasurface for Sub-6 Wireless Networks. In *20th USENIX Symposium on Networked Systems Design and Implementation* (Boston, MA, USA) *(NSDI '23)*. USENIX Association, Berkeley, CA, USA, 389–404. https://www.usenix.org/conference/nsdi23/presentation/li-xinyi

[34] Xinyi Li, Gaoteng Zhao, Ling Chen, Xinyu Zhang, and Ju Ren. 2024. RFMagus: Programming the Radio Environment With Networked Metasurfaces. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking* (Washington D.C., DC, USA) *(MobiCom '24)*. Association for Computing Machinery, New York, NY, USA, 16–30. doi:10.1145/3636534.3649344

[35] Zhuqi Li, Yaxiong Xie, Longfei Shangguan, Rotman Ivan Zelaya, Jeremy Gummeson, Wenjun Hu, and Kyle Jamieson. 2019. Towards Programming the Radio Environment with Large Arrays of Inexpensive Antennas. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI '19)*. USENIX Association, Berkeley, CA, USA, 285–300. https://www.usenix.org/conference/nsdi19/presentation/lizhuqi

[36] Christos Liaskos, Kostas Katsalis, Joan Triay, and Stefan Schmid. 2023. Resource management for programmable metasurfaces: Concept, prospects and challenges. *IEEE Commun. Mag.* 61, 11 (2023), 208–214. doi:10.1109/MCOM.002.2300012

[37] Gang Liu, F Richard Yu, Hong Ji, Victor CM Leung, and Xi Li. 2015. In-Band Full-Duplex Relaying: A Survey, Research Issues and Challenges. *IEEE Commun. Surveys Tuts.* 17, 2 (2015), 500–524. doi:10.1109/COMST.2015.2394324

[38] David G Luenberger. 1997. *Optimization by Vector Space Methods.* John Wiley & Sons, New York, NY, USA.

[39] Ruichun Ma, Lili Qiu, and Wenjun Hu. 2024. SurfOS: Towards an Operating System for Programmable Radio Environments. In *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks* (Irvine, CA, USA) *(HotNets '24)*. Association for Computing Machinery, New York, NY, USA, 132–141. doi:10.1145/3696348.

3696861
[40] Ruichun Ma, R. Ivan Zelaya, and Wenjun Hu. 2023. 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* (Madrid, Spain) *(MobiCom '23)*. Association for Computing Machinery, New York, NY, USA, Article 18, 15 pages. doi:10.1145/3570361.3592520
[41] 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 *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking* (Washington D.C., DC, USA) *(MobiCom '24)*. Association for Computing Machinery, New York, NY, USA, 62–76. doi:10.1145/3636534.3649347
[42] MG Chemicals. 2025. *841WB – EMF Paint for Electronics*. Retrieved June 26, 2026 from https://mgchemicals.com/products/conductive-paint/emf-paint/
[43] MG Chemicals. 2025. *842WBU – EMF Blocking Paint for Electronics*. Retrieved June 26, 2026 from https://mgchemicals.com/products/conductive-paint/conductive-water-based-paints/emf-blocking-paint/
[44] Mikrotik. 2023. *Cube 60Pro AC: A powerful point-to-point 60 GHz 802.11ay CPE for bridging distances up to 2.4 km*. Mikrotik. Retrieved June 26, 2026 from https://mikrotik.com/product/cube_60pro_ac
[45] Payam Nayeri, Fan Yang, and Atef Z Elsherbeni. 2018. *Reflectarray Antennas: Theory, Designs, and Applications*. John Wiley & Sons, Hoboken, NJ, USA.
[46] George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. 1978. An analysis of approximations for maximizing submodular set functions—I. *Math. Program.* 14, 1 (Dec. 1978), 265–294. doi:10.1007/BF01588971
[47] Robin Neuder, Marc Späth, Martin Schüßler, and Alejandro Jiménez-Sáez. 2024. Architecture for sub-100 ms liquid crystal reconfigurable intelligent surface based on defected delay lines. *Commun. Eng.* 3, 1 (2024), 70. doi:10.1038/s44172-024-00214-3
[48] Thomas Nitsche, Guillermo Bielsa, Irene Tejado, Adrian Loch, and Joerg Widmer. 2015. Boon and bane of 60 GHz networks: practical insights into beamforming, interference, and frame level operation. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies* (Heidelberg, Germany) *(CoNEXT '15)*. Association for Computing Machinery, New York, NY, USA, Article 17, 13 pages. doi:10.1145/2716281.2836102
[49] Nikolaos Nomikos, Themistoklis Charalambous, Ioannis Krikidis, Dimitrios N Skoutas, Demosthenes Vouyioukas, Mikael Johansson, and Charalabos Skianis. 2015. A Survey on Buffer-Aided Relay Selection. *IEEE Commun. Surveys Tuts.* 18, 2 (2015), 1073–1097. doi:10.1109/COMST.2015.2512441
[50] OnGo Alliance. 2022. *OnGo Private 5G Deployment Guide*. Technical Report. OnGo Alliance. Retrieved June 26, 2026 from https://ongoalliance.org/wp-content/uploads/2022/05/OnGo-Private-5G-Deployment-Guide-1.0.0.pdf
[51] Anurag Pallaprolu, Winston Hurst, Sophia Paul, and Yasamin Mostofi. 2023. I Beg to Diffract: RF Field Programming With Edges. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking* (Madrid, Spain) *(MobiCom '23)*. Association for Computing Machinery, New York, NY, USA, Article 52, 15 pages. doi:10.1145/3570361.3613266
[52] Hao Pan, Lili Qiu, Bei Ouyang, Shicheng Zheng, Yongzhao Zhang, Yi-Chao Chen, and Guangtao Xue. 2023. PMSat: Optimizing Passive Metasurface for Low Earth Orbit Satellite Communication. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking* (Madrid, Spain) *(MobiCom '23)*. Association for Computing Machinery, New York, NY, USA, Article 43, 15 pages. doi:10.1145/3570361.3613257
[53] Kun Qian, Lulu Yao, Xinyu Zhang, and Tse Nga Ng. 2022. MilliMirror: 3D printed reflecting surface for millimeter-wave coverage expansion. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking* (Sydney, NSW, Australia) *(MobiCom '22)*. Association for Computing Machinery, New York, NY, USA, 15–28. doi:10.1145/3495243.3517024
[54] Qualcomm. 2022. *Vision, market drivers, and research directions on the path to 6G*. White Paper. Qualcomm. Retrieved June 26, 2026 from https://www.qualcomm.com/content/dam/qcomm-martech/dm-assets/documents/Qualcomm-Whitepaper-Vision-market-drivers-and-research-directions-on-the-path-to-6G.pdf
[55] Theodore S Rappaport, James N Murdock, and Felix Gutierrez. 2011. State of the Art in 60-GHz Integrated Circuits and Systems for Wireless Communications. *Proc. IEEE* 99, 8 (2011), 1390–1436. doi:10.1109/JPROC.2011.2143650
[56] Theodore S Rappaport, Shu Sun, Rimma Mayzus, Hang Zhao, Yaniv Azar, Kevin Wang, George N Wong, Jocelyn K Schulz, Mathew Samimi, and Felix Gutierrez. 2013. Millimeter Wave Mobile Communications for 5G Cellular: It Will Work! *IEEE Access* 1 (2013), 335–349. doi:10.1109/ACCESS.2013.2260813
[57] Real Digital. 2024. *RFSoC 4x2*. Real Digital. Retrieved June 26, 2026 from https://www.realdigital.org/hardware/rfsoc-4x2
[58] Kenneth Rose. 1998. Deterministic Annealing for Clustering, Compression, Classification, Regression, and Related Optimization Problems. *Proc. IEEE* 86, 11 (Nov. 1998), 2210–2239. doi:10.1109/5.726788
[59] Julia Schwarzbeck, Robin Neuder, Marc Späth, and Alejandro Jiménez-Sáez. 2026. Scalable mm-Wave Liquid Crystal Reconfigurable Intelligent Surfaces based on the Delay Line Architecture. arXiv:2601.11307 [eess.SP] https://arxiv.org/abs/2601.11307
[60] Claude E Shannon. 1948. A Mathematical Theory of Communication. *Bell Labs Tech. J.* 27, 3 (July 1948), 379–423. doi:10.1002/j.1538-7305.1948.tb01338.x
[61] Khushboo Singh, Muhammad U Afzal, Maria Kovaleva, and Karu P Esselle. 2020. Controlling the Most Significant Grating Lobes in Two-Dimensional Beam-Steering Systems With Phase-Gradient Metasurfaces. *IEEE Trans. Antennas Propag.* 68, 3 (March 2020), 1389–1401. doi:10.1109/TAP.2019.2955244
[62] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. 2016. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. In *13th USENIX Symposium on Networked Systems Design and Implementation* (Santa Clara, CA, USA) *(NSDI '16)*. USENIX Association, Berkeley, CA, USA, 193–206. https://www.usenix.org/conference/nsdi16/technical-sessions/presentation/sur
[63] Verizon. 2020. *Professional Services Wi-Fi 6 RF Survey and Design Statement of Work to Verizon Professional Services Service Attachment*. Retrieved June 26, 2026 from https://www.verizon.com/business/service_guide/reg/cp_wi-fi_6_rf_survey_and_design_SOW.pdf
[64] Teng Wei and Xinyu Zhang. 2017. 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* (Snowbird, Utah, USA) *(MobiCom '17)*. Association for Computing Machinery, New York, NY, USA, 42–55. doi:10.1145/3117811.3117832
[65] Timothy Woodford, Kun Qian, and Xinyu Zhang. 2024. Metasight: High-Resolution NLoS Radar with Efficient Metasurface Encoding. In *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems* (Istanbul, Turkiye) *(SenSys '23)*. Association for Computing Machinery, New York, NY, USA, 308–321. doi:10.1145/3625687.3625803
[66] Shuanshuan Wu, Rachad Atat, Nicholas Mastronarde, and Lingjia Liu. 2017. Improving the Coverage and Spectral Efficiency of Millimeter-Wave Cellular Networks Using Device-to-Device Relays. *IEEE Trans. Commun.* 66, 5 (2017), 2251–2265. doi:10.1109/TCOMM.2017.2787990
[67] Renjie Zhao, Timothy Woodford, Teng Wei, Kun Qian, and Xinyu Zhang. 2020. M-Cube: a millimeter-wave massive MIMO software radio. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking* (London, United Kingdom) *(MobiCom '20)*. Association for Computing Machinery, New York, NY, USA, Article 15, 14 pages. doi:10.1145/3372224.3380892
[68] Kai Zheng, Wuqiong Zhao, Timothy Woodford, Renjie Zhao, Xinyu Zhang, and Yingbo Hua. 2024. Enhancing mmWave Radar Sensing Using a Phased-MIMO Architecture. In *Proceedings of the 22nd Annual International Conference on Mobile Systems, Applications and Services* (Minato-ku, Tokyo, Japan) *(MOBISYS '24)*. Association for Computing Machinery, New York, NY, USA, 56–69. doi:10.1145/3643832.3661865
[69] Kai Zheng, Wuqiong Zhao, and Xinyu Zhang. 2026. MetaPolar: A Low-Cost, Passive, mmWave Radar-Readable Metasurface Road Sign. In *Proceedings of the 2026 ACM/IEEE International Conference on Embedded Artificial Intelligence and Sensing Systems* (Saint Malo, France) *(SenSys '26)*. Association for Computing Machinery, New York, NY, USA, 200–213. doi:10.1145/3774906.3800464
[70] Anfu Zhou, Shaoqing Xu, Song Wang, Jingqi Huang, Shaoyuan Yang, Teng Wei, Xinyu Zhang, and Huadong Ma. 2019. Robot Navigation in Radio Beam Space: Leveraging Robotic Intelligence for Seamless mmWave Network Coverage. In *Proceedings of the Twentieth ACM International Symposium on Mobile Ad Hoc Networking and Computing* (Catania, Italy) *(MobiHoc '19)*. Association for Computing Machinery, New York, NY, USA, 161–170. doi:10.1145/3323679.3326514

Appendices are supporting material that has not been peer-reviewed.

A Notation

Table 1 summarizes the key notation used in the paper.

Table 1: Table of notation.

Symbol	Description
$\mathcal{M}$, $\mathcal{R}$, $\mathcal{K}$	Sets of MTS locations, Rx locations, angular grid
$S_i[k]$	Relative signal power at angular direction k of MTS i
$\mathbf{\Psi} = [\alpha_m, \mathbf{p}_m]$	Optimization parameters (placement & patterns)
$\mathcal{U}(S)$	Sum-rate objective function
SNR_r	Signal-to-noise ratio (SNR) at receiver r
$W_{i\to j}[k]$	Directional gain from emitter i to target j at angular bin k
η_j	Efficiency of MTS j
$\Delta_i[k]$	Directional score (= $R_i[k] + D_i[k]$) at MTS i, bin k
$R_i[k]$, $D_i[k]$	Relay (MTS-facing) & direct (user-facing) score components
$S_{d,r}$, $S_{d,t}$	Receiver and transmitter diversity scores
N_s, N_b	Number of receiver sectors and transmitter beams
$\bar{r}_i$, $\bar{r}_b$	Normalized sector/beam rates ($\sum_i \bar{r}_i = 1$, $\sum_b \bar{r}_b = 1$)
$\text{RSS}_{i(k)\to j}$	RSS from AP i with beam k to user j
$I_{i(k)\to j}$	Aggregate interference at user j from AP i, beam k
$\alpha_{\{t,r\}}$	Diversity exponent ($\in (0,1]$; lower → more diversity)
$y_i(S)$	Outgoing power at node i under pattern S
$W[m,r]$	Routing weight from MTS m to user r (APRO)
$x_{r,m}$	Signal contribution from MTS m to user r
β	Temperature parameter for APRO annealing

B Proof for Major–Minor Flow Topology

The following proofs assume a 2D angular space for simplicity, which can be easily extended to 3D.

B.1 Definitions and Assumptions

For notational simplicity without loss of generality, we also assume the same size for all MTSs. Define the angular grid $\mathcal{K}$, which is proportional to the MTS size. To easily represent the beam resolution limit, let $|\mathcal{K}|$ equal the number of MTS unit cells. For any emitter $i \in \{0\} \cup \mathcal{M}$ and any target $j \in \mathcal{M} \cup \mathcal{R}$, let $W_{i\to t}[k] \in [0,1)$ denote the directional gain from i to t at angular grid $k \in \mathcal{K}$, which involves the channel loss and Tx/MTS element pattern gain, and follows

$$\sum_{t\in\mathcal{M}\cup\mathcal{R}} W_{i\to t}[k] \le 1, \quad \forall i \in \{0\} \cup \mathcal{M}. \tag{9}$$

Specially, for $i = 0$ (the Tx), $W_{0\to t}[k] \triangleq \sup_{b\in\mathcal{B}_0} W^{(b)}_{0\to t}[k]$ is the Tx beam pattern envelope over codebook $\mathcal{B}_0$, a simplified model of the directional Tx. Following the principle in § 4.1, we assume coherent combination within each MTS, but non-coherent power combination across MTSs. We further relax the MTS beam pattern as a power constraint over the angular grid. The incoming power at MTS j is

$$y_j = \eta_j \sum_{i\in\{0\}\cup\mathcal{M}} y_i \sum_{k\in\mathcal{K}} W_{i\to j}[k] S_i[k], \quad \forall j \in \mathcal{M}, \tag{10}$$

where η_j is the efficiency of MTS j. Define the linear operator $T(S) \in \mathbb{R}_+^{|\mathcal{M}|\times|\mathcal{M}|}$ by

$$T_{j,i}(S) = \eta_j \sum_{k\in\mathcal{K}} W_{i\to j}[k] S_i[k]. \tag{11}$$

Then, Eq. (10) can be written as $y_{\mathcal{M}} = T(S) y_{\mathcal{M}} + b$ with $b \triangleq (\eta_j \sum W_{0\to j} S_0)_{j\in\mathcal{M}}$. Thus we have

$$y(S) = (I - T(S))^{-1} e_0, \tag{12}$$

where e_0 ejects the Tx's power. Eq. (12) holds with Lemma B.1.

Lemma B.1 (Stability via Passivity). *Let $\eta_{\max} \triangleq \max_j \eta_j < 1$. Then $\|T(S)\|_1 \le \eta_{\max} < 1$, and thus the spectral radius $\rho(T(S)) < 1$. Consequently, $(I - T(S))^{-1}$ exists, is non-negative, and*

$$y(S) = \sum_{l=0}^{\infty} T(S)^l e_0. \tag{13}$$

Proof. Using Eq. (9) and the constraints of $S[k]$ in Eq. (1), $\sum_j T_{j,i} = \sum_j \eta_j \sum W_{i\to j} S_i \le \eta_{\max} \sum S_i \le \eta_{\max}$, so $\|T(S)\|_1 \le \eta_{\max} < 1$. Positivity and the Neumann series [19] follow. □

Using power combination across MTSs, the received SNR at $r \in \mathcal{R}$ is

$$\text{SNR}_r(S) = \kappa \sum_{i\in\{0\}\cup\mathcal{M}} y_i(S) \sum_{k\in\mathcal{K}} W_{i\to r}[k] S_i[k], \tag{14}$$

where $\kappa \triangleq P_0/N_0$, and P_0 and N_0 are the constant Tx and noise power. The first-order optimality is based on the optimization problem in Eq. (1), with recursion in Eq. (12).

B.2 First-Order Optimality

Proof of Lemma 4.1. The user marginal slope

$$\Phi_r \triangleq \frac{\partial \mathcal{U}}{\partial p_r} = \frac{\kappa}{1 + \text{SNR}_r} > 0, \tag{15}$$

where p_r is the received power at r. Define vector ϕ with

$$\phi_j \triangleq \sum_{r\in\mathcal{R}} \Phi_r \sum_{k\in\mathcal{K}} W_{j\to r}[k] S_j[k], \quad \forall j \in \mathcal{M}. \tag{16}$$

Let the adjoint (downstream value) vector $\Lambda \in \mathbb{R}_+^{|\mathcal{M}|}$ solve

$$\Lambda = (I - T(S)^\mathsf{T})^{-1} \phi. \tag{17}$$

Since $\rho(T) < 1$ from Lemma B.1, Eq. (17) has a unique non-negative solution. Gateaux derivative (directional score) yields Eq. (2), where

$$D_i[k] \triangleq \sum_{r\in\mathcal{R}} \Phi_r W_{i\to r}[k], \quad R_i[k] \triangleq \sum_{j\in\mathcal{M}} \lambda_j \eta_j W_{i\to j}[k]. \tag{18}$$

Using the Karush–Kuhn–Tucker (KKT) stationarity [32], let $\mu_i \ge 0$ be the multiplier of the constraint $\sum_k S_i[k] \le 1$ in Eq. (1), and $\nu_i[k]$ be the multiplier of the non-negativity constraint $S_i[k] \ge 0$. Stationarity and complementary slackness give

$$\begin{cases} S_i[k] > 0 \implies \Delta_i[k] = \sup_{k'} \Delta_i[k'], \\ S_i[k] = 0 \implies \Delta_i[k] < \sup_{k'} \Delta_i[k']. \end{cases} \tag{19}$$

Define the optimal slopes and scores at $S^\star$:

$$\begin{cases} \Phi_r^\star = \dfrac{\kappa}{1 + \mathrm{SNR}_r(S^\star)}, \\ \Lambda^\star = (I - T(S^\star)^\mathsf{T})^{-1}\phi^\star, \\ \Delta_i^\star[k] = R_i^\star[k] + D_i^\star[k]. \end{cases} \tag{20}$$

Let the maximizer set be $\mathcal{A}_i \triangleq \arg\max_k \Delta_i^\star[k]$. Using the KKT support characterization at optimum $S^\star$, we have, if $S^\star[k] > 0$ then $\Delta_i^\star[k] = \sup_{k'} \Delta_i^\star[k']$. Equivalently, $\mathrm{supp}(S_i^\star) \subseteq \mathcal{A}_i^\star$. □

B.3 Networked Metasurfaces Topology

Angular separability condition G_0. As introduced in the main text, we formalize the angular separability condition. In many room-and-corridor layouts, and for sufficiently fine angular grids $\mathcal{K}$, each bin k for $W_{i\to t}[k]$ has a dominant "first interaction" class: it primarily points to *(i)* a boundary patch supporting a metasurface hop, or *(ii)* open-space user directions. This motivates the following separability condition, which is exact in the idealized case and approximately true in practice. For each emitter $i \in \{0\} \cup \mathcal{M}$ and angular bin $k \in \mathcal{K}$,

$$\Big(\sum_{j\in\mathcal{M}} W_{i\to j}[k]\Big) \cdot \Big(\sum_{r\in\mathcal{R}} W_{i\to r}[k]\Big) = 0. \tag{21}$$

Equivalently, each bin k is either MTS-facing ($W_{i\to j}[k] > 0$ for some $j \in \mathcal{M}$ but $W_{i\to r}[k] = 0$ for all $r \in \mathcal{R}$), or receiver-facing ($W_{i\to r}[k] > 0$ for some $r \in \mathcal{R}$ but $W_{i\to j}[k] = 0$ for all $j \in \mathcal{M}$), (or neither, in which case both sums are zero).

Proof of Proposition 4.2. Fix any node $i \in \{0\} \cup \mathcal{M}$ and consider an optimizer $S^\star$ of **(P)**. By Lemma 4.1, the KKT support characterization gives

$$S_i^\star[k] > 0 \implies \Delta_i^\star[k] = \sup_{k'\in\mathcal{K}} \Delta_i^\star[k']. \tag{22}$$

Under G_0, each bin is either MTS-facing or user-facing. Hence, for any k:

$$\begin{aligned} &(\text{MTS-facing}) \Rightarrow D_i^\star[k] = 0 \text{ and } \Delta_i^\star[k] = R_i^\star[k], \\ &(\text{user-facing}) \Rightarrow R_i^\star[k] = 0 \text{ and } \Delta_i^\star[k] = D_i^\star[k]. \end{aligned} \tag{23}$$

Therefore, if the major condition Eq. (3) holds at node i, *i.e.*, $\max_k \Delta_i^\star[k] > \max_k D_i^\star[k]$, then no user-facing bin can maximize $\Delta_i^\star[\cdot]$ since on those bins $\Delta_i^\star[k] = D_i^\star[k] \le \max_k D_i^\star[k] < \max_k \Delta_i^\star[k]$. By Eq. (22), $S_i^\star$ can only place positive mass on metasurface-facing bins, *i.e.*, node i acts as a *major-flow* node (relaying to other metasurfaces).

Similarly, if $\max_k D_i^\star[k] > \max_k R_i^\star[k]$ at node i, then no metasurface-facing bin can maximize $\Delta_i^\star[\cdot]$ since on those bins $\Delta_i^\star[k] = R_i^\star[k] \le \max_k R_i^\star[k] < \max_k D_i^\star[k] \le \max_k \Delta_i^\star[k]$. Again by Eq. (22), $S_i^\star$ allocates only to user-facing bins, so node i acts as a *minor-flow* node (fan-out coverage).

Combining the above cases, under exact G_0 any optimizer $S^\star$ of **(P)** induces a major–minor structure in the following sense: if $\max_k \Delta_i^\star[k] > \max_k D_i^\star[k]$ (Eq. (3)), then node i allocates only to MTS-facing bins (major-flow behavior); if $\max_k D_i^\star[k] > \max_k R_i^\star[k]$, then node i allocates only to receiver-facing bins (minor-flow behavior). If $\max_k D_i^\star[k] = \max_k R_i^\star[k]$, then both types of bins may lie in $\arg\max_k \Delta_i^\star[k]$, and the optimizer may mix them without contradicting the major/minor interpretation.

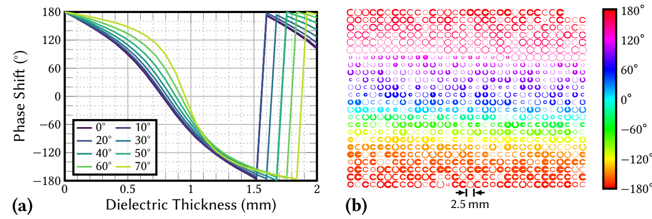

Figure 18: Unit cell phase profiles of two designs simulated using HFSS. (a) 3D-printed (wavy ground plane) [53]. (b) SRR-based [41].

In typical room/corridor indoor geometries with strong blockage, upstream metasurfaces often have limited visibility to uncovered user regions, making $\max_k D_i^\star[k]$ small; moreover, RSS/rate saturation at high SNR reduces the marginal slope $\Phi_r = \kappa/(1 + \mathrm{SNR}_r)$ for already-strong users, further suppressing $D_i[k]$. In contrast, well-placed downstream metasurfaces can have strong geometric hops (large $W_{i\to j}[k]$ for some k), and the downstream adjoint value λ_j aggregates the marginal utility of reaching many users through subsequent passive hops, making $\max_k R_i^\star[k]$ large. Hence, the score ordering that triggers major vs. minor behavior is typical in practice, yielding the major–minor flow topology described in the proposition.

Near-optimality under approximate G_0. In practice, limited angular resolution and scattering may cause a bin k to weakly couple to both metasurfaces and users. We model this by an ε-overlap: for all (i, k),

$$\min\Big\{\sum_{j\in\mathcal{M}} W_{i\to j}[k],\ \sum_{r\in\mathcal{R}} W_{i\to r}[k]\Big\} \le \varepsilon, \qquad \varepsilon \ll 1. \tag{24}$$

Then each bin still has a dominant first-interaction class, and the non-dominant term contributes only an $O(\varepsilon)$ perturbation to $\Delta_i^\star[k] = R_i^\star[k] + D_i^\star[k]$ (scaled by the finite multipliers $\lambda^\star, \Phi^\star$). Consequently, the KKT maximizer set of $\Delta_i^\star[\cdot]$ is unchanged whenever the relay-vs-direct score gap has a margin larger than $O(\varepsilon)$; otherwise both choices are within $O(\varepsilon)$ of optimal, which justifies the near-optimal statement in typical indoor environments. □

In summary, the major–minor topology is effective because major flows create high-gain relay chains that transport energy around geometric blockages, while minor flows provide angular diversity for broad coverage—aligned with FlowForm's coverage expansion objective.

C Detailed Metasurface Designs

C.1 Unit Cell Models

3D-printed metasurface [53]. Following the transmission line model in [53] for the metal-backed dielectric cuboid and considering oblique incidence at angle θ_i, the reflection coefficient is:

$$\Gamma = e^{-j2\arctan\left(\frac{1}{R\sqrt{\epsilon_r}}\tan\left(\sqrt{\epsilon_r}k_0 t_d\right)\right)}, \tag{25}$$

where $k_0 = 2\pi/\lambda$ is the free-space wavenumber, ϵ_r is the dielectric permittivity, t_d is the dielectric thickness, and $R = (1-\sin^2\phi_i \sin^2\theta_i)/\cos\theta_i$ is an angle-dependent coefficient. The phase profile for different θ_i simulated in HFSS (Fig. 18(a)) matches the model in Eq. (25). The dielectric constant of the Bambu PLA Basic filament [5] is measured as 2.57 at 60 GHz.

SRR-based metasurface [41]. The phase shift with the updated design (shown in Fig. 6(d)) is simulated using HFSS, where the opening of the split ring can be either 0° or 90°, with an approximate phase shift of 180° between them.

Alternative flat-back 3D-printed metasurfaces [65]. Unlike the wavy ground in [53], [65] has a flat ground, where each unit cell is a combination of dielectric and air. This makes each unit cell thicker, increasing the coupling effect (observed in Fig. 8(b)).

C.2 Metasurface Fabrication

FlowForm metasurface's phased array can be fabricated with both 3D printing [53, 65] and hot-stamping [41] techniques, as shown in Fig. 6(a). Our technique improves upon [53] by substituting copper ground with silver ground for better resonance/reflection efficiency, and enhances [41]'s durability and stability by laminating the hot-stamped SRR metasurface with a Polyethylene Terephthalate (PET) film. In both fabrication methods, unit cells are half-wavelength sized and spaced (2.5 mm at 60 GHz). We choose [53]'s design over [65], with FlowForm's frontal view being completely flat, because it offers a thinner profile, higher efficiency, and a completely flat frontal surface. The back is painted with water based conductive silver paint for common ground. Once the water-based paint dries, it leaves FlowForm metasurface with an ultra low resistance (*i.e.*, less than 1 Ω) ground layer on the uneven side. While other conductive paints (e.g., copper [10, 53], nickel [42]) are also applicable, silver offers superior electrical conductance that minimizes reflection loss. We also identify [10] as a cost-effective alternative. Each 80 × 80 surface costs about $2 USD for the 3D-printed metasurface (about 3 h printing with 45 g PLA filament [5] for $0.5 USD, and conductive paint [43] for $1.5 USD), and about $0.5 USD for the SRR-based metasurface.

C.3 Focal Distance Optimization

Algorithm 2 Joint Focal Distance Optimization

Require: Quantized heights $\{h_1, \dots, h_N\}$ per MTS, initial focal distances $\{d_1, \dots, d_M\}$
Ensure: Optimized focal distances $\{d_1^*, \dots, d_M^*\}$
Initialize focal coefficient $c_m \leftarrow 1/(2d_m)$, LR $\eta \leftarrow 0.5$
for $t = 1$ to $T_{\max}$ (*e.g.*, 25) **do**
  Apply quadratic phase correction $\Delta\Phi_n = k_0 \|r_n\|^2 c_m$ for each element n of MTS m
  Compute SNR(c) via channel model
  Update $\mathbf{c} \leftarrow \mathbf{c} + \eta \cdot \nabla_{\mathbf{c}}\text{rate}(\mathbf{c})$ via auto differentiation
end for
return $d_m^* = 1/(2c_m)$ for each MTS m

C.4 Late Acceptance Hill-Climbing for Refining Metasurface Unit-Cells

Algorithm 3 LAHC for Discrete Height Optimization

Require: Phase patterns $\{\boldsymbol{\phi}_1, \dots, \boldsymbol{\phi}_M\}$ for M MTSs
Ensure: Optimized discrete heights $\{\mathbf{h}_1, \dots, \mathbf{h}_M\}$
Initialize $h_{m,i} \leftarrow$ quantized height to $\phi_{m,i}$ for each MTS m, element i
Initialize objective history list L of size K (*e.g.*, 25)
$L \leftarrow [J(\mathbf{h})] \times K$ w/ $J(\mathbf{h}) \triangleq \text{rate}(\mathbf{h}) - \lambda \cdot \text{roughness}(\mathbf{h})$
for $t = 1$ to N_{eval} (*e.g.*, 1000) **do**
  Select random MTS m & random patch of elements
  Propose $h'_{m,i} \leftarrow h_{m,i} \pm 1$ discrete level (*e.g.*, 0.08 mm, in wrap-around) for each element i in patch
  Evaluate J_{new} globally with modified heights
  if $J_{\text{new}} > L[t \bmod K]$ **or** $J_{\text{new}} > J_{\text{current}}$ **then**
    Accept: $h_{m,i} \leftarrow h'_{m,i}$, $J_{\text{current}} \leftarrow J_{\text{new}}$
  end if
  Update history: $L[t \bmod K] \leftarrow J_{\text{current}}$
end for
return $\{\mathbf{h}_1, \dots, \mathbf{h}_M\}$

Neighborhood structures. We use patch-based moves with variable neighborhood search (VNS): patch sizes escalate from small to large as the search stagnates, combined with iterated local search (ILS) perturbations to escape local optima.

D Additional Experiment Results

All the experiment scenes are shown in Table 2, and photos of all measurement scenes are shown in Fig. 19.

Table 2: All experiment scenes.

Scene	Area (m^2)	Description
A	68.6	L-shaped hallway connecting to a study area
B	133.0	A study area with multiple hallways
C	324.6	T-shaped common area with a long corridor
D	382.9	A multi-arm lobby area with a zigzag shape
E	462.0	A hall connected to corridors with multi arms

Coverage expansion results are shown in Fig. 20 across 5 different scenes.

Figure 19: Photos of all 5 scenes with FlowForm metasurfaces deployed.

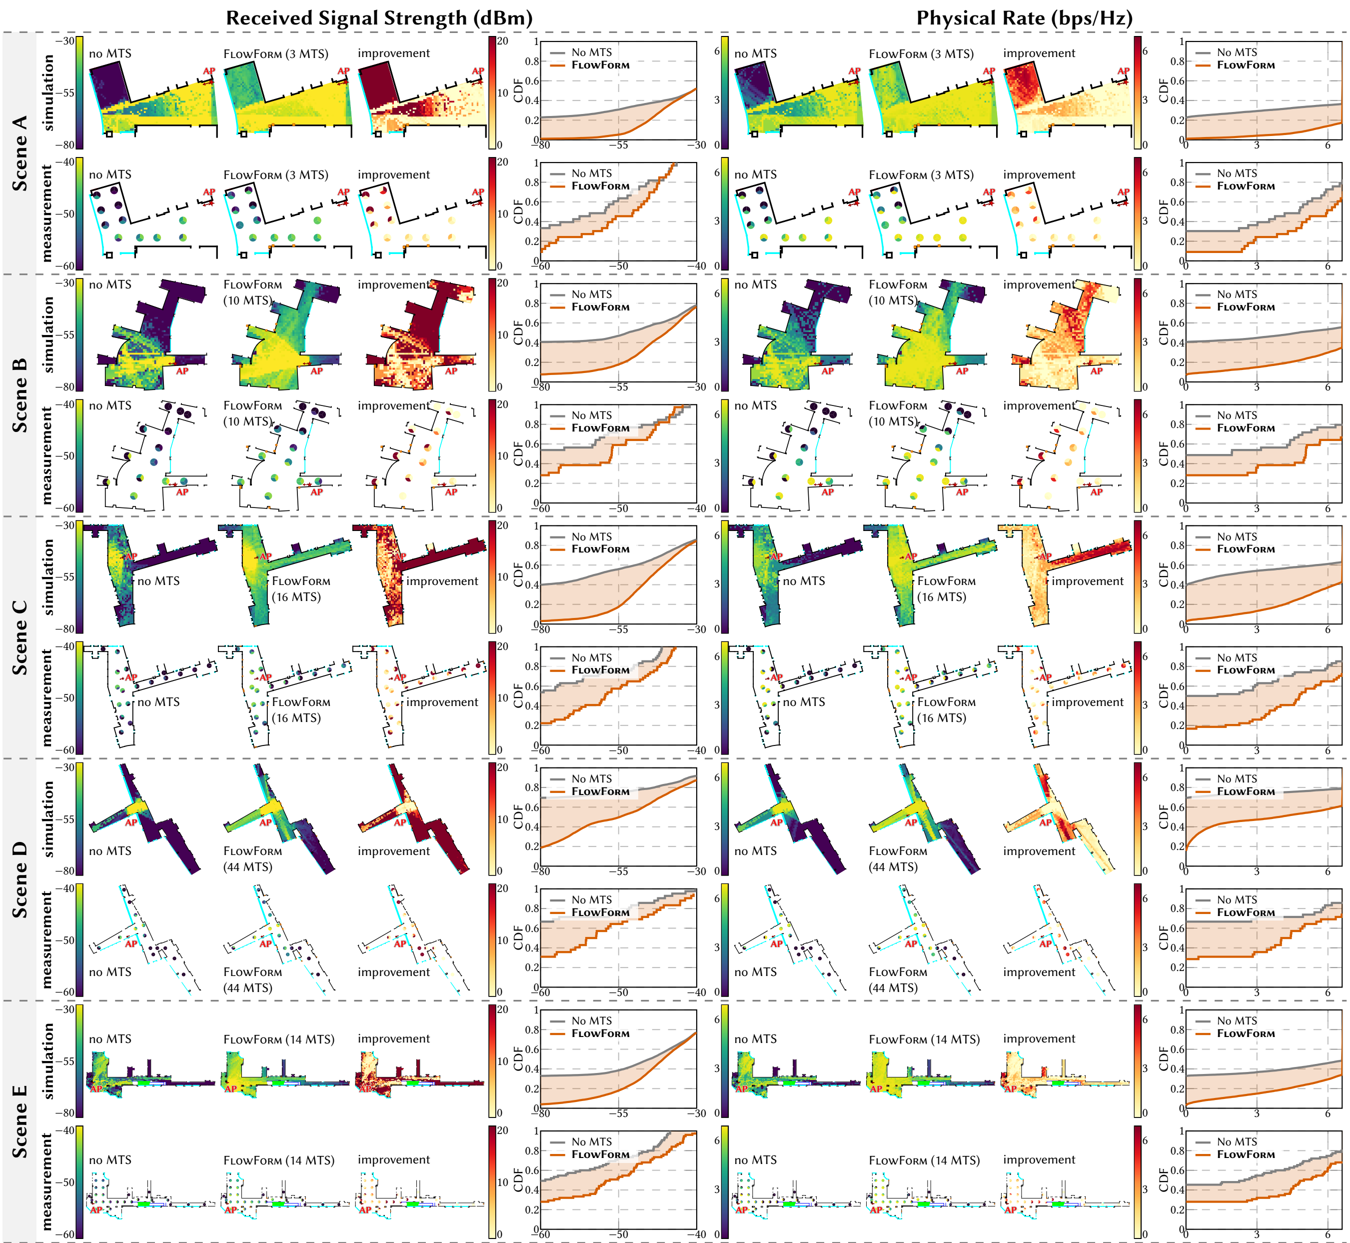

Figure 20: Experimental (simulation and measurement) results of coverage expansion.