

Customizing Indoor Wireless Coverage via 3D-Fabricated Reflectors

Xi Xiong, Justin Chan¹, Ethan Yu, Nisha Kumari, Ardan Amiri Sani², Changxi Zheng³, Xia Zhou
Dartmouth College, ¹University of Washington, ²UC Irvine, ³Columbia University

ABSTRACT

Judicious control of indoor wireless coverage is crucial in built environments. It enhances signal reception, reduces harmful interference, and raises the barrier for malicious attackers. Existing methods are either costly, vulnerable to attacks, or hard to configure. We present a low-cost, secure, and easy-to-configure approach that uses an easily-accessible, 3D-fabricated reflector to customize wireless coverage. With input on coarse-grained environment setting and preferred coverage (e.g., areas with signals to be strengthened or weakened), the system computes an optimized reflector shape tailored to the given environment. The user simply 3D prints the reflector and places it around a Wi-Fi access point to realize the target coverage. We conduct experiments to examine the efficacy and limits of optimized reflectors in different indoor settings. Results show that optimized reflectors coexist with a variety of Wi-Fi APs and correctly weaken or enhance signals in target areas by up to 10 or 6 dB, resulting to throughput changes by up to -63.3% or 55.1%.

CCS CONCEPTS

• **Networks** → **Network range**; *Wireless access networks*;

KEYWORDS

3D printing, wireless coverage, reflector

ACM Reference format:

Xi Xiong, Justin Chan¹, Ethan Yu, Nisha Kumari, Ardan Amiri Sani², Changxi Zheng³, Xia Zhou. 2017. Customizing Indoor Wireless Coverage via 3D-Fabricated Reflectors. In *Proceedings of BuildSys '17, Delft, Netherlands, November 8–9, 2017*, 10 pages. <https://doi.org/10.1145/3137133.3137148>

1 INTRODUCTION

Customizing the coverage of wireless networks explicitly for built environments is critical. By regulating the physical coverage of each wireless access point (AP), we can enhance signal reception in desired regions while weakening signals in others. It improves the efficiency of wireless infrastructure in buildings by mitigating the impact of building's insulations, partitions, and interior layouts, reducing harmful interference, and enhancing system security. Such physical confinement of wireless signals serves as a complimentary method to existing network security measures, such as encryption, and hence raises the barrier for attackers.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

BuildSys '17, November 8–9, 2017, Delft, Netherlands

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5544-5/17/11...\$15.00

<https://doi.org/10.1145/3137133.3137148>

Achieving this goal is particularly challenging indoors, because of the complex interactions of radio signals with the environment. Existing approaches rely on directional antennas to concentrate signals in desired directions. These approaches, however, have three shortcomings. *First*, they face a tradeoff between cost and control flexibility. Low-cost directional antennas [8] concentrate signals in a static direction, offering very few coverage shapes. Antenna arrays with more sophisticated control (e.g., arbitrary beam patterns, dynamic configuration) are costly, e.g., \$5000+ for Phocus array [6], \$200 for an 802.11ac AP [1]. It is because forming a narrow beam needs a large number of antennas, each with a separate RF chain [44]. These additional RF chains lead to a prohibitive cost. *Second*, they often fail to provide strong security guarantees. On the one hand, fixed-beam directional antennas cannot physically limit signals within an arbitrary area of interest due to their fixed radiation patterns. On the other hand, multi-user beamforming systems (e.g., 802.11ac APs) cannot differentiate intended and malicious clients and strengthen signals for both types of clients. *Third*, they require significant configuring efforts due to rich multi-path effects indoors [8, 14, 34]. Users must try different antenna configurations to identify the one best matching the desired coverage.

In this paper, we present a low-cost, physically-secure, and easy-to-configure approach that customizes each AP's coverage without requiring directional antennas. Given the increasing popularity and easy accessibility of 3D printers¹, we study the use of 3D-fabricated reflectors produced by 3D digital manufacturing ("3D printing") to control signal propagation in the space. Specifically, we place a signal reflector around an AP, where the reflector shape is computationally optimized, taking into account the environment (e.g., interior layouts, partitions, AP locations) and the desired signal distribution (i.e., target areas to strengthen or weaken signal strength). The reflector reflects wireless signals to realize a desired coverage. Indeed, anecdotal experiments [2] have demonstrated substantial (29.1%–57.2%) bandwidth gain by placing a soda can behind a Wi-Fi AP to strengthen signal in one direction. Our work generalizes this idea by presenting a systematic approach to optimizing reflector shapes for enabling a rich set of signal distributions.

Our approach provides four benefits. First, it provides strong *physical security* by limiting the physical reach of wireless signals, hence creating a virtual wall for wireless signals. Second, it relies on a low-cost (\$35), reproducible 3D reflector, which can be easily replaced upon substantial changes in the environment or coverage requirement. Third, it offers an easily-accessible and easy-to-configure solution to non-expert users. Users only need to specify coverage requirements and a coarse environment model², with which our system computes a reflector shape tailored to the built environment. Finally, it is applicable to commodity low-end Wi-Fi

¹Many online printing services [4, 7] are available when 3D printers are absent.

²Indoor models can be obtained using full-fledged systems (e.g., Google Tango, Microsoft HoloLens) or other 3D geometry reconstruction techniques [17, 25, 40].

	Physical Security	Config. Effort	Hardware Cost	Directionality Gain
Fixed-beam DA	Limited	High	Low	Good
Configurable-beam DA	Yes	High	High	Best
Multi-user beamforming	No	Low	High	Best
Our solution	Yes	Low	Low	Good

Table 1: Comparing our solution to directional antennas (DA) or beamforming.

APs without directional or multiple antennas. Table 1 compares our solution to its alternatives.

Our work comprises two technical components: *First*, we design an effective optimization algorithm that optimizes reflector 3D shapes for a target wireless coverage. During this process, we represent a reflector shape as a parametric model [39] in computer graphics to ensure surface smoothness and in turn the feasibility of 3D fabrication. The shape optimization leverages a 3D wireless modeling to evaluate the effectiveness of a candidate shape and improve the shape iteratively. We guide a Simulated Annealing algorithm [28] using local gradient descent to sample the shape space more efficiently. We also extend our optimization to deal with multiple APs and jointly optimize their reflector shapes. *Second*, we develop an efficient modeling approach that uses 3D ray tracing to simulate radio signal propagation and signal’s interaction with objects in a 3D environment. We consider signal’s reflection, transmission, and diffraction through objects. For APs with multiple antennas, we also take into account the antenna location and orientation to trace radio signals accurately in the 3D space.

We 3D print optimized reflectors, test them with various Wi-Fi APs (including the latest 802.11ac AP) via signal and throughput measurements in two indoor settings. Our findings are as follows:

- Optimized reflectors correctly adjust signal distribution towards the target coverage. Resulting signal strength can decrease by up to 10 dB and increase by 6 dB, leading to throughput differences from -63.3% to 55.1%;
- Optimized reflectors coexist nicely with various Wi-Fi APs including MIMO APs (e.g., TP-Link AP and Netgear R7000). They allow multiple APs to collaboratively serve a region, or to confine each AP’s coverage to enhance security and reduce interference;
- The optimized reflector is relatively easy to place. Its efficacy is not sensitive to slight placement offsets, tolerating up to 10° offset in orientation and 10 cm offset in the distance to the AP.
- Given an environment model, our system computes an optimized reflector shape in 23 minutes on a laptop (2.2. GHz Intel Core i7).

Our work presents a low-cost, secure, and easy-to-configure approach that can coexist with directional communication via multi-antenna beamforming. Physical reflectors regulate signal distribution across regions at the macro level, while beamforming enables finer-grained signal enhancement at individual clients within the coverage. While we use Wi-Fi as an example, our approach can be generalized to other wireless bands (e.g., light, millimeter waves, acoustics). It can improve efficiency of diverse wireless networks in built environments, and potentially help in delivering higher-order services for buildings, e.g., occupancy detection, lighting control.

2 SYSTEM OVERVIEW

Our system takes three inputs: 1) a digitized, simplified environmental model including the main room structure (e.g., walls); 2)

information of wireless APs: their locations, the number of antennas, and positions of antennas if they are external antennas; and 3) the target wireless coverage specified as areas where users aim to strengthen or weaken the received signal strength.

The core of our system is an iterative stochastic optimization process that searches for a 3D reflector shape for each AP, so that collectively they achieve the target wireless coverage. Starting from an initial 3D shape and reflector position, our algorithm perturbs the current shape and estimates the resulting signal distribution with the new shape. Based on an objective function that measures the quality of a reflector shape, it then chooses to accept or reject the perturbed shape before moving on to the next iteration. Evaluating the objective function requires a 3D wireless propagation model that predicts the spatial distribution of received signal strength. Our wireless propagation model takes into account the indoor environment and simulates the radio waves interacting with environment objects. The process stabilizes at a final shape until no further improvement can be made to better match the target.

Finally, we output the shape and placement of the reflector, then the user fabricates the optimized reflector shape and coats it with a thin metal layer (e.g., aluminum foil) to enhance its ability to reflect wireless signals. The fabricated reflector is then mounted around the wireless AP to realize the customization of wireless coverage.

Next, we describe our optimization of reflection shape, followed by wireless propagation modeling and reflector fabrication.

3 REFLECTOR SHAPE OPTIMIZATION

We optimize the reflector shape, aiming to generate a target wireless coverage for a given environment. We first present parameterizing the 3D geometry of a reflector, followed by our shape optimization procedure, as well as its extension to deal with multiple APs.

3.1 Representing the Reflector Shape

To represent a 3D surface, we seek a shape parameterization that can express a large space of feasible shapes and yet entail a low control degrees of freedom for the sake of computational performance. A naive solution is to represent a shape as a triangle mesh and optimize the positions of mesh vertices. However, this approach has a large number of optimization variables (i.e., the positions of mesh vertices). It results into an optimization problem in a high dimensional space, which is rather computationally expensive. Furthermore, the resulting shape might not be well-formed and thus cannot be fabricated in practice.

We overcome this challenge by leveraging the Non-uniform rational B-spline (NURBS) surface [39], a shape parameterization commonly used in computer graphics and engineering design [18]. The NURBS surface offers appealing properties: 1) it is flexible, supporting a large variety of shapes including standard analytical shapes such as spheres and free-form shapes; 2) it guarantees a smooth surface, facilitating the fabrication; and 3) it enables fine-grained control with low complexity using a small number of control points.

The NURBS surface is defined by NURBS curves. A NURBS curve consists of a set of weighted control points, a knot vector and its order. We aim to fit control points on the curve, which then determine the curve shape. The knot vector is a list of ascending numbers, defining where and how control points affect the NURBS




Figure 1: NURBS surface [39] to generate and control a 3D surface using a few (4×4) control points.

curve. By manipulating the knot vector, we can decide whether the curve passes through or passes by certain control points. The size of the knot vector is equal to the number of control points plus the order. The order defines the number of nearby control points that influence any given point on the curve.

A NURBS surface is calculated as the tensor product of two NURBS curves. Thus, it has two parametric directions (u and v) and two corresponding orders and knot vectors. For our purposes, we predefine its knot vectors and orders and then manipulate its shape by changing positions of the control points. A NURBS surface $\Omega(u, v)$ is constructed as:

$$\Omega(u, v) = \sum_{i=1}^l \sum_{j=1}^w R_{i,j}(u, v) P_{i,j}, \quad (1)$$

with

$$R_{i,j}(u, v) = \frac{N_{i,n}(u)N_{j,m}(v)w_{i,j}}{\sum_{p=1}^l \sum_{q=1}^w N_{p,n}(u)N_{q,m}(v)w_{p,q}}$$

as rational basis functions, where we have $l \times w$ control points $P_{i,j}$, and $w_{i,j}$ is the corresponding weight. $N_{i,n}(u)$ is the i^{th} B-spline basis function of degree n [19]. In our model, a shape Ω is defined by $l \times w$ control points $P_{i,j}$. We set their weights as 1, which is a common usage. Figure 1 illustrates two example NURBS surfaces with 4×4 control points, where we lift the center control points to generate a concave shape on the left and then lift the four corner points to generate the shape on the right.

3.2 Optimizing the Reflector Shape

To search for the reflector shape optimized for a target coverage, we start with a flat plane as the reflector shape Ω and then perturb Ω over iterations. In each iteration, we evaluate the effectiveness of a candidate shape Ω by the objective function $F(\Omega)$, where $F(\Omega)$ measures the difference between the desired coverage and the coverage C_Ω resulting from reflector shape Ω . We estimate C_Ω by running the 3D wireless modeling described in Section 4. Specifically, we divide the environment into small cells in uniform size ($1 \text{ m} \times 1 \text{ m}$ in our implementation). We apply the 3D wireless propagation modeling to predict signal strength at each cell's center. Assuming M_+ and M_- denote the areas where users aim to strengthen and weaken signals respectively, we compute $F(\Omega)$ as:

$$F(\Omega) = \sum_{i \in M_+ \cup M_-} \|C_{Target}(i) - C_\Omega(i)\|^2, \quad (2)$$

where $C_\Omega(i)$ is the signal strength in dBm at cell i after placing reflector shape Ω , and $C_{Target}(i)$ is the target signal strength of cell i . To derive $C_{Target}(i)$, we first compute signal $C(i)$, which is the estimated signal strength at cell i when no reflector is placed. We then add/subtract δ , which is the expected signal enhancement/reduction




Figure 2: (a) shows the objective function $F(\Omega)$ as we perturb two control points of the NURBS surface. (b) illustrates our search method. In the $(k-1)$ -th iteration (Ω_{k-1}), we apply gradient descent to seek the local optimum Ω'_k and choose Ω'_k as the next candidate.

at cell i . To determine δ , we tested different materials on their performance of signal enhancement and attenuation at distances from 1 to 3 meters. Results (Figure 5) show that the signal change is at most 15 dB. Thus we set δ to 15 dB and write $C_{Target}(i)$ as:

$$C_{Target}(i) = \begin{cases} C(i) + \delta & \text{if } i \in M_+ \\ C(i) - \delta & \text{if } i \in M_- \end{cases}. \quad (3)$$

Here we use a cell-dependent target value $C_{Target}(i)$, rather than a uniform signal upper and lower bound, because $C_{Target}(i)$ defines an equal range (δ) above or below $C(i)$. As a result, a given amount of signal enhancement or reduction leads to the same amount of change in $F(\cdot)$, regardless of the cell location. This, however, is no longer guaranteed if a uniform bound is used, because of the quadratic nature of $F(\cdot)$. Thus, $C_{Target}(i)$ ensures the optimization process is unbiased across cells. Finally, the optimization process searches for $\Omega^* : \Omega^* = \operatorname{argmin}_\Omega F(\Omega)$, where Ω^* is the 3D reflector shape that leads to signal distribution best matching the target.

While the optimization problem appears standard, the search space is daunting and many local optimums exist (Figure 2(a)). Simple local search methods such as hill climbing can easily be stuck at local optimums. We need efficient algorithm to identify the global optimal. To achieve this goal, we consider simulated annealing (SA) [28], which allows the iterations to opportunistically escape from the current local search area even if the escape leads to an increase in the objective function. The escape likelihood p is determined by two parameters in the algorithm: current temperature T and the increase in the objective function. SA keeps examining candidate shapes until T reaches the minimal temperature (0). When a candidate shape Ω is examined, the current temperature is cooled at a rate r . SA accepts Ω if $F(\Omega)$ is lower than the previous candidate. Otherwise, it accepts Ω with a probability p . p is adapted over iterations. In the beginning when T is higher, p is also higher so SA tends to explore other areas in the shape space. As T decreases to 0, p approaches 0 so that it gradually settles at an optimum.

However, SA can require a fairly large number of iterations as it randomly samples the search space. To achieve better results, we propose to guide SA's iterations using gradient descent. The key idea is to consider characteristics of a local search area for determining the next candidate. Specifically, in each iteration of SA, instead of randomly generating a shape as the candidate, we apply gradient descent to seek the local minimum as the next candidate. Take the k th iteration as an example, we first generate a random shape Ω_k , then we calculate the shape gradient $\nabla F(\Omega_k)$ of objective

Algorithm 1 Shape Optimization

```

1: initialize  $\Omega_k, k = 0, T = T_{max}$ 
2: while  $T > T_{min}$  do
3:    $k \leftarrow k + 1$ 
4:    $F(\Omega_{k-1}) \leftarrow Eq.(2)$ 
5:    $\Omega_k \leftarrow perturb(\Omega_{k-1})$ 
6:   while do
7:      $\nabla F(\Omega_k) \leftarrow getGradient(F(\Omega_k))$ 
8:      $\lambda \leftarrow getStepSize()$ 
9:     if  $F(\Omega_k - \lambda \nabla F(\Omega_k)) < F(\Omega_k)$  then
10:       $\Omega_k \leftarrow \Omega_k - \lambda \nabla F(\Omega_k)$ 
11:     else
12:       break
13:     end if
14:   end while
15:    $F(\Omega_k) \leftarrow Eq.(2)$ 
16:    $p \leftarrow e^{\frac{F(\Omega_{k-1}) - F(\Omega_k)}{T}}$ 
17:   if  $F(\Omega_{k-1}) \geq F(\Omega_k)$  or  $rand[0, 1] \leq p$  then
18:      $\Omega^* \leftarrow \Omega_k$ 
19:   end if
20:    $T \leftarrow T \cdot r$ 
21: end while
22: return  $\Omega^*$ 

```

function $F(\Omega_k)$ at Ω_k :

$$\nabla F(\Omega_k) = \lim_{||d\Omega_k|| \rightarrow 0} \frac{F(\Omega_k + d\Omega_k) - F(\Omega_k)}{d\Omega_k}, \quad (4)$$

where $d\Omega_k$ is obtained by slightly changing the control points of Ω_k . Then we apply backtracking line search to find an appropriate step size λ . Finally we take $\Omega_k - \lambda \nabla F(\Omega_k)$ as Ω_k and repeat Eq. (4) until we find a local optimum Ω'_k . We take Ω'_k as the candidate instead of Ω_k to go over the accept/reject procedure. This method directly iterates from one local minimum to another (Figure 2(b)), and thus is more efficient to approach the global optimum than SA's random sampling [31]. Algorithm 1 lists the detail.

Accelerating the Search. Each iteration in our search can be time-consuming, because deriving $\nabla F(\Omega_k)$ requires altering the positions of Ω_k 's control points one by one. The running time of each iteration is linear with the number of Ω_k 's control points. To speed up the gradient computation, we leverage the simultaneous perturbation stochastic approximation (SPSA) algorithm [13, 46], which perturbs all parameters (i.e., control points) simultaneously with a random perturbation vector Δ to estimate the gradient of each parameter. Thus, it approximates the gradient computation using only two calculations of the objective function, regardless of the parameter dimension (i.e., the number of control points in our problem). Additionally, we run these iterations as parallel threads to further shorten the process. As a result, an iteration take 1.71 seconds on average on a MacBook Pro (2.2 GHz Intel Core i7).

3.3 Extending to Multiple APs

We now extend the above shape optimization to handle multiple APs. We classify APs into two types: 1) *Collaborative APs*: APs that are deployed by the same entity, e.g., a user or an enterprise deploying multiple APs in a home or workplace. These APs collaboratively serve a region to provide wireless coverage and user's device automatically connects to the AP with the strongest signal; 2) *Non-Collaborative APs*: APs that are deployed by different entities. Each AP serves users in its own pre-defined coverage region, without the knowledge nor any control of other APs.

For collaborative APs, we jointly optimize their reflector shapes so that the resulting signal coverage best matches the target. Here the signal coverage map is defined based on the strongest signal received at a location from all collaborative APs. Thus, let $\mathbb{O} = \{\Omega_1, \dots, \Omega_M\}$ denote a set of candidate reflector shapes for M collaborative APs, where Ω_j is the reflector shape for AP j . Then its objective function $F(\mathbb{O})$ is similar to Eq.(2):

$$F(\mathbb{O}) = \sum_{i \in M_t \cup M_r} \|C_{Target}(i) - C_{\mathbb{O}}(i)\|^2, \quad (5)$$

where $C_{\mathbb{O}}(i)$ is the signal strength at cell i after placing reflector shapes $\mathbb{O}$ at all APs. $C_{Target}(i)$ is computed following Eq.(3), but with $C(i)$ as the signal received at cell i from the strongest AP. Thus the search is to seek the optimal set of shapes $\mathbb{O}^* = \text{argmin}_{\mathbb{O}} F(\mathbb{O})$. To jointly optimize M reflectors, we interleave the perturbation of each AP's reflector shape across iterations. Specifically, in an iteration, we perturb only one AP's reflector shape to seek its next candidate shape while fixing other APs' reflector shapes, and then move on to perturbing the next AP's reflector shape in the next iteration. Each iteration of the optimization procedure is similar to Algorithm 1. We omit the algorithm details in the interest of space.

For non-collaborative APs, each AP has its own pre-defined coverage region without the knowledge of other APs, thus each AP's reflector shape is optimized separately following Algorithm 1. For non-collaborative APs operating on the same channel, they can interfere with one another if their coverage regions are nearby. The impact of interference can be minimized if the information (e.g., location, reflector shape) of other interfering APs is available to the shape optimization algorithm. We can estimate the interference at each cell and consider signal-to-interference ratio (SINR) as the target (Eq. (2)) for shape optimization. We leave it for future work.

4 EFFICIENT 3D WIRELESS MODELING

An essential part of our shape optimization is an efficient modeling that simulates wireless signal propagation in a given environment for evaluating the efficacy of a candidate reflector shape. Most existing models either fall far short in modeling accuracy [29, 36], or require expensive measurement or computation overhead. To achieve a better tradeoff, we choose 3D ray tracing for its best accuracy [26, 35, 53] and design schemes to speed up its computation. In particular, we choose the Shooting-and-Bouncing Ray (SBR) launching algorithm [21, 30, 43], which launches a number of rays from the transmitter and traces all their possible paths to reach a receiver. Next we describe the key steps of our modeling.

Ray Launching. We model each antenna of a Wi-Fi AP as a signal source³. To achieve accurate ray tracing, rays need to be uniformly spread from the transmitter, i.e., angles between any two adjacent rays are uniform. To do so, we use the geodesic sphere ray launching [51]. A geodesic sphere (Figure 3) is formed by tessellating faces of a regular polyhedron and extrapolating the intersection points on the surface of a sphere [27]. The geodesic vertices provide equivalent angular separation around the entire sphere [43]. Thus, we can place the transmitter at the sphere center and emit a ray to every geodesic vertex on the sphere to guarantee uniform

³More antennas imply more signal sources and thus higher computation overhead in wireless modeling. They also impose more physical constraints to place the reflector.




Figure 3: Geodesic sphere for uniform ray launching, where a ray is emitted from the sphere center to each geodesic vertex on the sphere. The sphere is generated by tessellating a regular icosahedron (a). (b) and (c) show the resulting spheres after tessellating each triangle surface in (a) into 4 and 256 triangles, respectively.

ray launching. In our implementation, we consider tessellating an icosahedron (20 triangular faces and 12 vertices) [54]. To achieve high resolution, we divide each edge into $N = 8$ segments and thus tessellate each triangle surface into $N^2 = 64$ triangles (Figure 3(c)). It leads to $10N^2 + 2 = 642$ launched rays (i.e., geodesic vertices) with the average radial angular separation θ of $\frac{1}{N} \sqrt{\frac{4\pi}{5\sqrt{3}}} \approx 0.1506$ rad, where the unit of 4π is steradian [21].

Ray Tracing. We track each ray's interaction with environmental objects. Given the wavelength of Wi-Fi signals, we consider three types of interactions (Figure 4(a)): *transmissions* (rays penetrate the objects), *reflections* (a ray is bounced over a smooth surface), and *diffractions* (a ray hitting an object edge is diffracted as a set of rays in a cone shape, Figure 4(b)). To model diffraction, we leverage the Uniform Theory of Diffraction (UTD) used by the fast-wave acoustics simulation in computer graphics [52]. It has also been applied in modeling RF propagation in buildings [37]. Similar to [24, 26], we do not consider wave phase as we average signal strengths in each $1 \text{ m} \times 1 \text{ m}$ cell when evaluating coverage. We do not model other wave phenomena like scattering since they have negligible impact on the resulting signal map [26]. We also do not consider the near-field effect, since we calculate signal strength at locations at least meters away from the antenna, while the Fraunhofer distance [11] defining the near field is tens of centimeters in our setting⁴.

Each type of ray interaction contributes to additional energy loss of a ray, in addition to its signal degradation over distance. To integrate all these contributors, we choose a partition model [9, 26] to calculate each ray's signal power at a receiver location. The model consists of four parts: 1) the signal degradation over distance, represented by the pathloss exponent α ; 2) the reflection attenuation, which is the product of the reflection coefficient β and the number of reflections; 3) the transmission attenuation, which is the product of the transmission coefficient γ and the number of times that a ray penetrates obstacles; and 4) the diffraction attenuation, which is the product of the diffraction coefficient λ and the number of diffractions. We do not consider multiple diffractions and reflection-diffraction [43].

Formally, let P_i^j denote the power in dBm contributed by the i^{th} ray of j^{th} AP after traveling a distance of d_i to reach a receiver, we can calculate P_i^j as:

$$P_i^j = P_0^j - 10\alpha \log_{10}(d_i/d_0) - \beta N_{i,ref} - \gamma N_{i,trans} - \lambda N_{i,diff} - \beta' N'_{i,ref} - \gamma' N'_{i,trans}, \quad (6)$$

⁴Fraunhofer distance is calculated as $2D^2/\lambda$, where D is the largest physical linear dimension of the antenna and λ is the wavelength.




Figure 4: Ray tracing.

where P_0^j (in dBm) is the reference power of j^{th} AP at distance d_0 , pathloss exponent α captures how quickly the signal degrades over distance, $N_{i,ref}$, $N_{i,trans}$, $N_{i,diff}$ are the number of reflections, transmissions, and diffractions that ray i experiences, respectively. Since the reflector surface is designed to be highly reflective, we consider different transmission and reflection coefficients (β' , γ') for the reflector. $N'_{i,ref}$ and $N'_{i,trans}$ are the number of times ray i penetrates and is reflected by the reflector, respectively.

We can adapt the model to various environments easily as parameters (P_0^j , α , β , γ , λ , β' , γ') are calibrated only once. The calibration needs measurements only at a few sampled locations, rather than a site survey. The best-fit parameter values are identified by applying simulated annealing [26] to minimize modeling errors. These parameters can then be reused at all locations of an environment.

Ray Reception. Finally to calculate the received signal strength at a receiver location, we sum the power of each ray within the reception zone Z of this receiver. Formally, let P^j (in dBm) be the received power from AP j , we have $P^j = 10 \log_{10} \sum_{i \in Z} 10^{P_i^j/10}$, where P_i^j is the power contributed by ray i from AP j calculated by Eq. (6). We consider the reception zone as a sphere with radius of $\theta d/\sqrt{3}$ centered at the receiver [42], where θ is the average radial angular separation between adjacent rays launched from the transmitter, and d is the length of a ray's propagation path to the receiver. The radius of the reception sphere considers the fact that rays are spread out as they propagate. Figure 4 shows the reception sphere for each ray, where ray2 travels the longest distance to reach receiver and thus has the largest reception sphere.

Speeding Up Ray Tracing. While offering higher accuracy, ray tracing incurs heavy computation, mainly because of calculating intersections of a large number of rays and triangle meshes⁵. To speed up the ray tracing process, we index the triangle meshes with a Kd-Tree [12]. We construct a Kd-Tree of the bounding boxes containing these triangle meshes. For each ray, we first search the tree to identify the bounding boxes that the ray intersects. Only for the triangle meshes in those bounding boxes, we examine the ray-triangle intersection, which avoids many unnecessary intersection tests. For each ray-box and ray-triangle intersection test, we apply prior algorithms [22, 32] to improve the efficiency.

5 REFLECTOR FABRICATION

With the optimized shape Ω^* , the last step is to realize it as a physical reflector. We represent Ω^* as its corresponding control points of the NURBS surface, calculate the coordinates of each mesh point based on Eq. (1), and store all mesh points in an .obj file. We

⁵In our implementation, the indoor model has 107 triangle meshes and the ray tracing launched 1926 rays, producing $\approx 10\text{K}$ rays after multiple reflections and transmissions.




Figure 5: Different plane reflector’s ability to reflect and attenuate Wi-Fi signals, at distances from 1 to 3 meters.

edit the .obj file in Blender, an open-source 3D computer graphics software. We add thickness and export it to a .stl file for a 3D printer (MakerBot) to print it. Its build volume is 25.2 cm × 20 cm × 15.0 cm. Overall fabricating a 20 cm × 20 cm reflector costs no more than \$35 (one large spool of MakerBot PLA Filament).

We add a thin layer of metal to the plastic reflector surface because metals have exceedingly high conductivities and thus are effective reflectors and attenuators of radio waves [3]. To determine the metal material, we systematically test three types of metal sheets made of silver, copper, and aluminum, which are all 0.1-mm thick. We also test aluminum sheet with 0.25-mm thickness to evaluate the impact of the metal-layer thickness. We attach each metal sheet to a 30 cm × 30 cm plastic (the same material used by 3D printers in fabrication) to form the final reflector. We place the reflector at the back of an AP to test its ability to reflect Wi-Fi signals to a receiver. To test its ability to attenuate signals, we place the reflector between the AP and the receiver. For both tests, we vary the distance of the receiver to the AP from 1 m to 3 m.

In Figure 5, we see that copper and aluminum perform similarly in enhancing and attenuating Wi-Fi signals, and silver performs slightly better. Increasing sheet thickness moderately improves its ability to attenuate signals, but not its reflection property. Given the lower cost of aluminum and the difficulty of applying a 0.1-mm sheet to an uneven reflector surface, we choose to cover the reflector surface with a 0.024-mm household aluminum foil.

We also observe that reflectors are better at weakening than strengthening Wi-Fi signals, mainly because of Wi-Fi’s wavelength and the energy loss when signals penetrate the reflector. When Wi-Fi signals interact with the reflector, the signal either penetrates, or is absorbed, or is reflected by the reflector. Only the reflected energy can be directed to strengthen the signal at locations before the reflector, while both the reflected and the absorbed energy contribute to the signal attenuation at locations behind the reflector.

6 EVALUATION

We evaluate our approach by testing optimized reflectors in two indoor scenes. We seek to understand its capability in affecting wireless signal distribution and throughput, implications on enhancing security and reducing interference, its sensitivity to reflector placement and size, and other performance microbenchmarks.

6.1 Experimental Setup

Indoor Scenes. We experiment two indoor scenarios with different room layouts: (a) *workspace scenario*: a 19 m × 13 m indoor area with a narrow (7.5-m long) hallway connecting a research lab and

two offices; (b) *home scenario*: a 16 m × 12 m area where a spacious lobby (5.8 m × 5.1 m) is surrounded by three rooms. Figure 6(a) and Figure 7(a) show their floor maps, respectively. Although rooms are furnished in both scenes, the 3D environment models used in our shape optimization contain only walls and doors (we will evaluate the impact of including furniture in environment models in § 6.5). Experiments are conducted during working hours with moving users around (walking, working at their desks, or standing and talking to others). For the simplicity and feasibility of the user input, we assume the same material for objects in the environment. To consider heterogeneous materials, we can adapt the reflection, transmission and diffraction coefficients in Eq. (6) based on materials, and track what materials a ray has traveled through.

Wi-Fi APs. We use Netgear R7000 (IEEE 802.11ac) as default APs. R7000 has three antennas and operates on both 2.4 GHz and 5 GHz frequency. We configure it to transmit at 10 mW and collect RSS values at 2.4 GHz band. We also test two other popular APs: Linksys WRT54GL (IEEE 802.11g) and TP-Link WR841N (IEEE 802.11n), both operating on 2.4 GHz frequency and equipped with two external antennas. To minimize interference from external APs, we analyze channel usage status using a mobile app (Wi-Fi Analyzer) and set the AP to operate on the least congested channel (channel 9 in our environment).

Ground Truth Signal Collection. Although our approach does not require exhaustive site survey measurements to compute reflector shapes, such measurements are necessary for us to evaluate the impact of optimized reflectors on signal distribution. To gather signal measurements, we divide each area into 1 m × 1 m cells and average received signal strength (RSS) values within each cell. We also sample a few locations (blue circles in Figure 6(a) and Figure 7(a)) to measure the throughput.

To automate RSS measurements in the 3D space, we apply a drone-based method in [55]. We program an AR Drone 2.0 to collect RSS at specified locations. We reuse drone’s built-in Wi-Fi radio to receive beacons from our AP and record RSS values. To decide the measurement duration per location, we let the drone hover over a location, measure for 1 minute and 10 seconds respectively, and compare the mean and standard deviation of RSS. Results reveal that 10 seconds are sufficient for collecting stable RSS statistics.

6.2 Efficacy of Optimized Reflector

We start by examining the overall efficacy of the optimized reflector in a single-AP setting. We consider a target coverage in Figure 6(a), where we mark areas with received signals to be strengthened by ticks and areas with signals to be weakened by crosses. We run our shape optimization algorithm to derive the optimized reflector shape, fabricate the reflector (20 cm × 20 cm in size), and place it around AP’s antennas. Figure 6(d) shows the optimized 3D shape, its estimated radiation pattern, and its placement. The radiation pattern is generated by simulating signal change at one meter away. For this target coverage, an antenna is behind the reflector and the others are in front of the reflector. We measure changes in the resulting RSS and throughput.

Impact on Signal Distribution. We collect the signal maps before and after placing the reflector around the AP. We also compare




Figure 6: Efficacy of an optimized reflector for achieving a target wireless coverage (a), where areas users aim to strengthen the signal are marked by green ticks and areas to weaken the signal are marked by red crosses. (b) and (c) show a reflector in a simple concave shape and its resulting signal change in dB, while (d), (e), and (f) show a reflector in optimized shape, and the resulting changes in signal distribution and throughput. The optimized reflector shape leads to a signal distribution better matching the target.

the results to that of a naive reflector shape (the concave shape used in anecdote experiments, Figure 6(b)). Figure 6(c) and (e) show the RSS change (in dB), where positive numbers indicate signal enhancements and negative numbers indicate signal declines. We observe that the optimized reflector correctly adjusts signals in all target areas, weakening signals in the hallway and room2 by 10 dB while strengthening signals in other target areas by 6 dB. It achieves the goal by blocking an antenna that emits signals to the hallway and room2 while reflecting the signals of the other two antenna towards room1. The naive reflector, however, generates a simple radiation pattern that uniformly weakens signals in all areas in the back of the reflector, and thus fails to meet requirements of all target areas. The result demonstrates the necessity and efficacy of our optimization, which considers the indoor layout to customize the reflector shape and enables more flexible control.

A side effect of the optimized reflector is that in order to weaken signals in the hallway and room2, it also slightly weakens the RSS in the right bottom of the lab. It is a sacrifice made by the shape optimization to reach an overall signal distribution better matching the target. As we further analyze signal change at individual cells, 83.1% of cells have their signals correctly strengthened or weakened, demonstrating the overall efficacy of the optimized reflector.

Impact on Throughput. We further examine how signal-level changes translate into throughput differences at clients. We sample a few locations (marked as blue circles in Figure 6(a)) and measure the throughput at those locations before and after placing the reflector. In particular, we associate two laptops (MacBook Pro) with our AP. We fix the location of a laptop, while a user holds the other laptop walking around within each location to measure the throughput. We instrument one laptop to transmit 500-MB data to the other using the iperf utility and collect throughput statistics. We repeat the experiment for 10 rounds.

Figure 6(f) shows the percentage of throughput change under different APs with optimized reflectors. We also include error bars covering 90% confidence intervals. Overall throughput increases/decreases by up to 22.1%/36.7% in target areas. The throughput improvement at location 8 is small because its RSS is low, requiring a larger signal enhancement to switch to higher data rates. The throughput at location 9 slightly decreases because its RSS is slightly weakened by the reflector (Figure 6(e)). Most other locations in the lab experience improved throughput. Overall, the optimized reflector correctly adjusts throughput for 11 out of 12 locations. The result also shows that our reflectors can coexist nicely with MIMO APs (i.e., TP-Link AP and Netgear R7000).

6.3 Multi-AP Settings

We now move on to scenarios with multiple APs. As follows, we consider two types of APs described in § 3.3.

Collaborative APs. To evaluate the efficacy of jointly optimized reflectors, we deploy two APs in scene 2 and jointly optimize their reflector shapes for a target coverage in Figure 7(a). We measure the signal map before and after placing reflectors, where the RSS at a location is the signal from the stronger AP. We plot the signal change in Figure 7(b). The two optimized reflectors successfully reduce signal strength in lobby by up to 10 dB and increase signals up to 6dB in room 3, 4 and 5⁶. Overall, 91% of cells have their signals correctly strengthened/weakened. To further examine the impact on resulting bandwidth, we sample 9 locations (blue circles in Figure 7(a)) and plot in Figure 7(c) the throughput change brought by optimized reflectors. For 7 out of 9 locations, the throughput changes from -63.3% to 55.1%. Location 13 and 18 experience little change because of their minor signal changes (Figure 7(b)). Overall

⁶The numbers are slightly different from that in Figure 5 since the result here is measured at further distances and also affected by the environment.




Figure 7: Efficacy of optimized reflectors of two collaborative APs to achieve a target wireless coverage. In (a), areas where users aim to strengthen or weaken signals are marked by green ticks and red crosses, respectively. (b) shows the map of resulting signal change in dB. (c) plots the throughput improvement at sampled locations (blue circles in (a)).




Figure 8: Experiments on efficacy of reflectors confining wireless coverage and reducing interference. (a) shows the placement and desired Wi-Fi regions of two APs. (b) plots the packet loss rate change outside the desired region of each AP when optimized reflectors are placed. (c) presents reflectors' impact on SINR inside the AP's desired region.

the signal and throughput changes are notable for the majority of locations, demonstrating the efficacy of our joint optimization.

Non-Collaborative APs. For non-collaborative APs, we aim to confine each AP's coverage. We examine the implications on security/privacy and interference reduction with optimized reflectors.

1) Implications on Security/Privacy: We set up two Netgear R7000 APs in scene 1 and mark each AP's coverage region in Figure 8(a). We then fabricate reflectors for these APs to confine their coverage. We measure packet loss rates at sampled locations (marked as blue circles) before and after placing reflectors. We associate two laptops (MacBook Pro) with the same AP and instrument a laptop to transmit packets to the other using the iperf utility. Similarly to [45], we desire packet loss rates above 30% for locations outside an AP's coverage region, as many TCP and UDP based applications [10, 41] require a packet loss rate below 25%.

Figure 8(b) plots packet loss rates when a receiver outside an AP's region attempts to connect to this AP. Before placing any reflector, all locations have access to both APs. After placing optimized reflectors, we find that room2 (location 1 and 2) is unable to access the network of AP1 (50 - 60% packet loss rates), and locations at the lab (location 6, 7 and 8) can not access AP2. This demonstrates that optimized reflectors help confine Wi-Fi signal strength in the desired region. We also observe that for some locations outside an AP's region, packet loss rate does not change much after placing the reflector, e.g., the packet loss rate of AP2 at location 9 is nearly 0% with or without the reflector. It is because location 9 is relatively close to AP2 and receives a strong signal from AP2 even with the reflector. The same holds for location 3, 4, and 5. To address this limitation, we plan in future work to include transmit power as another parameter in our model to control coverage more precisely.

2) Reducing Interference: To quantify the benefit on interference reduction, we examine Signal to Interference Noise Ratio




Figure 9: Sensitivity to reflector placement offset.

(SINR) at locations in each AP's coverage region. We reuse the setting in Figure 8(a) and compute the SINR at all locations. Figure 8(c) shows that reflectors boost the SINR for most locations by up to 13 dB. The SINR at location 9 barely changes because while AP2's reflector weakens the interference, AP1's reflector also weakens the signal at this location, resulting into little change in its SINR.

6.4 Reflector Placement and Size

In addition to the reflector shape, the reflector placement and size can also affect reflector's efficacy. Next we examine system's sensitivity to these factors. In following experiments, we use the optimized reflector in Figure 6(d) as an example.

Placement Offset. We first study the impact of orientation offset, which can occur during manual placement. In the experiment, the reflector (20 cm × 20 cm) faces room1 and room2 to enhance signals, which corresponds to the 0° orientation offset. We rotate the reflector from -30° to 30° in a counterclockwise manner with 5° interval. We then plot the average signal enhancement of interested locations in room1 and room2 in Figure 9(a). During the experiment, we observe some large signal variation at certain locations (location 10) when the orientation offset is within 10°, but the variation is only within 0.51 dB for room1 and 0.83 dB for room2, indicating that the reflector can tolerate small orientation offset. Orientation offset larger than 10° results in slow decrease in signal enhancement.




Figure 10: Accuracy of 3D wireless propagation modeling, in comparison to the 2D modeling in prior work [16].

We then examine how the distance between the reflector and AP affects the performance. We fix the reflector’s orientation, vary its distance from 0 cm to 25 cm with 5-cm interval, and plot the average signal enhancement at room1 and room2 in Figure 9(b). We observe that offsets within 10 cm have negligible impact. Once the distance is above 10 cm, the signal enhancement starts to decrease. It is because as signals travel, they are further spread out with attenuated strength. Thus, a reflector further away from the AP reflects fewer signals and is less effective in affecting signal distribution.

Reflector Size. We also study the impact of reflector size on reflector’s ability to adjust signal distribution. A larger reflector reflects more signals and can be more effective, however its fabrication is harder and more costly. We aim to seek a proper size to achieve a good tradeoff. We fabricate the reflector in two sizes: 20 cm × 20 cm and 40 cm × 40 cm. Given that the height and width of all our APs are roughly 20 cm, a 20 cm × 20 cm reflector can cover all antennas. We place each reflector around the AP and measure the resulting RSS change at room1 and room2. As we compare the mean RSS change and the standard deviation for these two sizes, we observe negligible difference, indicating as long as the reflector covers all antennas, it needs not to be larger. This observation can be a guideline for determining reflector size for APs in other sizes.

6.5 Microbenchmarks

Finally, we examine the accuracy of our 3D wireless modeling and overall running time of the shape optimization.

Accuracy of 3D Wireless Modeling. Using scene 1 as an example, Figure 10 (a) and (b) plot CDFs of absolute RSS errors of our 3D wireless model (described in § 4), using measured RSS as ground truth. Figure 10 (a) is the result without any reflector, while Figure 10 (b) is for a concave shaped reflector placed around the AP. Overall the mean RSS error of 3D wireless modeling is 3 dBm.

Furthermore, we compare our 3D modeling to the 2D modeling in a prior work [16]. We observe that 3D modeling notably lowers the tail of RSS errors. The maximal error of 2D modeling is 11 dBm and 16 dBm, with and without reflector respectively, while they are 8 dBm and 12 dBm for 3D modeling. 3D modeling outperforms 2D modeling because of more accurate characterization of signal interaction in the third dimension. We also include the results when 2D/3D wireless modeling uses a finer-grained environment model that includes main furniture (e.g., desks, sofas). We observe that for both 2D and 3D modeling, adding furniture leads to negligible difference in resulting accuracy in both scenarios. The results indicate that coarse-grained 3D environmental models are sufficient.

Running Time. We run our shape optimization algorithm on a MacBook Pro (2.2 GHz Intel Core i7) and record the time to generate

an optimized shape for various target coverage requirements. We observe that the algorithm stabilizes after 23 minutes on average.

We further look into effectiveness of our speedup schemes: the Kd-tree used to speed up wireless modeling, the SPSA algorithm and multiple threads to speed up the search for optimized shape. Table 2 lists the running time of the optimization process when using either no speedup scheme, or only one speedup scheme, or all of them. Having 107 triangle meshes and 10K rays in the ray tracing model, we see that the Kd-tree accelerate around 1.6 times. For a complex model that has thousands of triangle meshes, the speedup of Kd-tree can be up to tens of times. Theoretically SPSA can reduce tens of time to calculate the gradient when we have 3×5 control points, however it only speeds up roughly 1.5 times, because the search for appropriate step size is another bottleneck. Multiple threads accelerate nearly 4 times. Together, they reduce the average running time from 207 minutes to 23 minutes.

Scheme	No speedup	SPSA only	Kd-tree only	Threads only	Full speedup
Time (minutes)	207	141	128	48	23

Table 2: Running time of our shape optimization with and without speedup schemes (SPSA, Kd-Tree, and multiple threads).

7 RELATED WORK

Configuring Wireless Coverage. Prior works have optimized AP placement to improve signal reception in certain areas [35]. However, moving the AP to enhance one area would result into the decline of signal strength in other areas. Thus, such methods are constrained in its flexibility. Another method is to use directional APs to confine wireless coverage to a specified region [45]. However, this method needs multiple costly directional APs. Our approach works with a single AP without directional antennas.

As for the use of reflectors, recent work [20, 49, 50] has studied reusing walls to reflect radio waves and control signal propagation. This method, however, relies on smart walls made of special materials and requires infrastructure-level changes. Similarly, a latest work [23] examined placing multiple metal plates in the environment to enhance wireless performance of a single AP. This approach also requires nontrivial changes in the environment by installing reflectors likely in large sizes. In contrast, our approach uses only a small reflector at the AP, is applicable in any environment, and supports multiple APs. Additionally, [23] optimizes reflector locations, whereas our work optimizes reflector shape. Another prior work [16] studied the feasibility of applying fabricated reflector to control wireless coverage. We advance this prior work in multiple fronts: a more sophisticated shape model, a more efficient shape optimization and extension to multiple APs, 3D wireless modeling, and extensive indoor experiments with optimized reflectors.

Directional Antennas. Directional antennas increase signal gain in a chosen direction and thus improve spatial reuse [8, 14, 15, 33, 38, 47, 48]. Researchers have studied steerable-beam directional antenna’s link quality outdoors [15] and its directionality indoors [14]. Its directionality greatly decreases due to rich multipath effects [8, 14, 34]. In comparison, we consider the influence of indoor layout when optimizing reflector shape. Besides, low-cost directional antennas (e.g., microstrip antennas [8], sectorized antennas [48]) only provide limited simple patterns. Our reflector is low-cost and flexible in radiation patterns.

Directional antennas often rely on multi-antenna beamforming, which electronically adjusts an array of omni-antennas to form narrow beams and maximize the signal strength for one or multiple users [5, 44]. Despite its superior performance, it has three main limitations. First, it cannot provide physical security, as it only enhances client's signal reception and cannot take into account client's location information (i.e., whether the client is inside the authorized location or not). As a result, it can end up forming a beam towards an attacker outside the authorized area. Second, it is more costly than our solution as it requires multiple antennas, each with a separate RF chain. Finally, it imposes expensive measurement overhead by collecting real-time Channel State Information (CSI).

8 CONCLUSION AND FUTURE WORK

We studied the design of a low-cost, 3D-fabricated reflector to customize wireless coverage. We demonstrated the efficacy of optimized reflectors with reflector prototypes and indoor experiments.

We summarize the limitations of our study and plans for future work. *First*, our current reflectors are in static shapes. Thus, upon substantial changes in the environment (e.g., removal/addition of walls), the reflector shape needs to be re-calculated and fabricated. Minor furniture changes have minimal impact on the resulting coverage, as shown in Figure 10, thus not requiring to replace the reflector. In the future, we will study reflectors made of trans-formable materials, enabling the reflector to automatically adapt its shape upon major changes. *Second*, the wavelength of Wi-Fi signals limits reflector's ability to reflect and block signals (Figure 5). Thus, changes on RSS and throughput have been moderate. Moving forward, we will examine higher frequency bands such as millimeter waves and visible light, where reflectors can block and reflect signals more effectively and cause more significant change in signal distribution. *Finally*, we will explore augmenting directional antennas with optimized reflectors and analyze their interplay.

REFERENCES

- [1] <http://www.dlink.com/uk/en/products/dir-890l-ac3200-ultra-wifi-router>.
- [2] <https://youtu.be/yz4aPaebe-k>.
- [3] http://stakeholders.ofcom.org.uk/binaries/research/technology-research/2014/building-materials-propagation/Building_Materials_and_Propagation.pdf.
- [4] 3D Hubs. <https://www.3dhubs.com/>.
- [5] 802.11ac In-Depth. White Paper, Aruba Networks.
- [6] Phocus array. <http://www.fidelity-comtech.com/products/phocus-array/>.
- [7] Shapeways. <https://www.shapeways.com/>.
- [8] AMIRI SANI, A., ZHONG, L., AND SABHARWAL, A. Directional antenna diversity for mobile devices: Characterizations and solutions. In *Proc. of MobiCom* (2010).
- [9] BAHL, P., AND PADMANABHAN, V. N. RADAR: An in-building RF-based user location and tracking system. In *Proc. of INFOCOM* (2000).
- [10] BALAKRISHNAN, H., ET AL. A Comparison of Mechanisms for Improving TCP Performance over Wireless Links. *IEEE/ACM Trans. Netw.* 5, 6 (1997), 756–769.
- [11] BALANIS, C. A. *Antenna theory: analysis and design*. John Wiley & Sons, 2016.
- [12] BENTLEY, J. L. Multidimensional binary search trees used for associative searching. *Communications of ACM* 18, 9 (Sept. 1975), 509–517.
- [13] BHATNAGAR, S., PRASAD, H. L., AND PRASHANTH, L. A. *Stochastic Recursive Algorithms for Optimization: Simultaneous Perturbation Methods*. Springer, 2013.
- [14] BLANCO, M., ET AL. On the effectiveness of switched beam antennas in indoor environments. In *Proc. of PAM* (2008).
- [15] BUETTNER, M., ET AL. A phased array antenna testbed for evaluating directionality in wireless networks. In *Proc. of MobiEval* (2007).
- [16] CHAN, J., ZHENG, C., AND ZHOU, X. 3D Printing Your Wireless Coverage. In *Proc. of HotWireless* (2015).
- [17] CHEN, J., BAUTEMBACH, D., AND IZADI, S. Scalable real-time volumetric surface reconstruction. In *Proc. of SIGGRAPH* (2013).
- [18] COTTELL, J. A., HUGHES, T. J., AND BAZILEVS, Y. *Isogeometric analysis: toward integration of CAD and FEA*. John Wiley & Sons, 2009.
- [19] DE BOOR, C., DE BOOR, C., DE BOOR, C., AND DE BOOR, C. *A practical guide to splines*, vol. 27. Springer-Verlag New York, 1978.
- [20] DUPRE, M., ET AL. Recycling radio waves with smart walls. In *International Conference on Metamaterials, Photonic Crystals and Plasmonics* (2015).
- [21] DURGIN, G., PATWARI, N., AND RAPPAPORT, T. S. An advanced 3D ray launching method for wireless propagation prediction. In *Proc. of VTC* (1997).
- [22] GLASSNER, A. S. *An introduction to ray tracing*. Elsevier, 1989.
- [23] HAN, S., AND SHIN, K. G. Enhancing Wireless Performance Using Reflectors. In *Proc. of INFOCOM* (2017).
- [24] HASSAN-ALI, M., AND PAHLAVAN, K. A new statistical model for site-specific indoor radio propagation prediction based on geometric optics and geometric probability. *IEEE transactions on Wireless Communications* 1, 1 (2002), 112–124.
- [25] IZADI, S., ET AL. KinectFusion: Real-time 3D Reconstruction and Interaction Using a Moving Depth Camera. In *Proc. of UIST* (2011).
- [26] JI, Y., BIAZ, S., PANDEY, S., AND AGRAWAL, P. ARIADNE: a dynamic indoor signal map construction and localization system. In *Proc. of MobiSys* (2006).
- [27] KENNER, H. *Geodesic math and how to use it*. Univ. of California Press, 1976.
- [28] KIRKPATRICK, S. Optimization by simulated annealing: Quantitative studies. *Journal of statistical physics* 34, 5-6 (1984), 975–986.
- [29] KOTZ, D., ET AL. Experimental evaluation of wireless simulation assumptions. In *Proc. of MSWiM* (2004).
- [30] LING, H., ET AL. Shooting and bouncing rays: Calculating the rcs of an arbitrarily shaped cavity. *IEEE Transactions on Antennas and Propagation* 37, 2 (1989), 194–205.
- [31] MARTIN, O. C., AND OTTO, S. W. Combining simulated annealing with local search heuristics. *Annals of Operations Research* 63, 1 (1996), 57–75.
- [32] MÖLLER, T., AND TRUMBORE, B. Fast, minimum storage ray/triangle intersection. In *ACM SIGGRAPH 2005 Courses* (2005), ACM, p. 7.
- [33] NAVDA, V., ET AL. MobiSteer: using steerable beam directional antenna for vehicular network access. In *Proc. of MobiSys* (2007).
- [34] NICULESCU, D., AND NATH, B. Vor base stations for indoor 802.11 positioning. In *Proc. of MobiCom* (2004), ACM.
- [35] PANJWANI, M. A., ET AL. Interactive computation of coverage regions for wireless communication in multifloored indoor environments. *Selected Areas in Communications, IEEE Journal on* 14, 3 (1996), 420–430.
- [36] PHILLIPS, C., SICKER, D., AND GRUNWALD, D. Bounding the error of path loss models. In *Proc. of DySPAN* (2011).
- [37] RAJKUMAR, A., ET AL. Predicting RF coverage in large environments using ray-beam tracing and partitioning tree represented geometry. *Wireless Networks* 2, 2 (1996), 143–154.
- [38] RAMACHANDRAN, K., ET AL. R2D2: Regulating Beam Shape and Rate As Directionality Meets Diversity. In *Proc. of MobiSys* (2009).
- [39] ROGERS, D. F. *An introduction to NURBS: with historical perspective*. Elsevier, 2000.
- [40] SANKAR, A., AND SEITZ, S. Capturing indoor scenes with smartphones. In *Proc. of UIST* (2012).
- [41] SAT, B., AND WAH, B. W. Analysis and evaluation of the Skype and Google-talk VoIP systems. In *IEEE International Conference on Multimedia and Expo*. (2006).
- [42] SCHAUBACH, K. R., ET AL. A ray tracing method for predicting path loss and delay spread in microcellular environments. In *Proc. of VTC* (1992).
- [43] SEIDEL, S. Y., AND RAPPAPORT, T. S. Site-specific propagation prediction for wireless in-building personal communication system design. *Vehicular Technology, IEEE Transactions on* 43, 4 (1994), 879–891.
- [44] SHEPARD, C., ET AL. Argos: Practical many-antenna base stations. In *Proc. of MobiCom* (2012).
- [45] SHETH, A., SESHAN, S., AND WETHERALL, D. Geo-fencing: Confining Wi-Fi coverage to physical boundaries. In *Pervasive Computing*. Springer, 2009, pp. 274–290.
- [46] SPALL, J. C. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control* 37, 3 (1992), 332–341.
- [47] SUBRAMANIAN, A. P., ET AL. A measurement study of inter-vehicular communication using steerable beam directional antenna. In *Proc. of VANET* (2008).
- [48] SUBRAMANIAN, A. P., ET AL. Experimental characterization of sectorized antennas in dense 802.11 wireless mesh networks. In *Proc. of MobiHoc* (2009).
- [49] SUBRT, L., AND PECHAC, P. Controlling propagation environments using intelligent walls. In *European Conference on Antennas and Propagation* (2012).
- [50] SUBRT, L., AND PECHAC, P. Intelligent walls as autonomous parts of smart indoor environments. *IET Communications* 6, 8 (May 2012), 1004–1010.
- [51] TAN, S., AND TAN, H. Modelling and measurements of channel impulse response for an indoor wireless communication system. In *Microwaves, Antennas and Propagation, IEE Proceedings* (1995), vol. 142, IET, p. 405.
- [52] TSINGOS, N., ET AL. Modeling acoustics in virtual environments using the uniform theory of diffraction. In *Proc. of SIGGRAPH* (2001).
- [53] VALENZUELA, R. A. A ray tracing approach to predicting indoor wireless transmission. In *Proc. of VTC* (1993).
- [54] WENNINGER, M. J. *Spherical models*, vol. 3. Courier Corporation, 1979.
- [55] YU, E., XIONG, X., AND ZHOU, X. Automating 3D Wireless Measurements with Drones. In *Proc. of WiNTECH* (2016).