Wireless Channel Lab: A Validated Browser-Based Ray Tracer and Link-Level Channel Simulator

Sean Kim
Department of Electrical and Computer Engineering
Princeton University, Princeton, NJ 08544, USA

***Abstract*—Site-specific ray tracers predict how radio energy moves through a building, and link-level simulators predict what the resulting channel does to a transmitted signal, but the two are rarely connected in one tool. This paper presents Wireless Channel Lab, an open-source simulator that runs entirely in a web browser and carries a single radio link from geometry to performance. A 2.5D propagation engine finds image-method reflections, ITU-R P.2040 slab reflection and transmission, knife-edge diffraction at wall ends, and a ground-reflected twin for every ray, and a beam-tracing pass reduces the third-order image tree by a factor of 34 without changing any result. The resulting channel, including its average SNR, Rician K-factor, delay spread, ray taps, and departure angles, is passed directly to modules for fading, OFDM, modulation, diversity, beamforming, and MIMO. Every mechanism is verified against closed-form theory using reference code written independently in Python. Deterministic results agree to within 1.1×10^{-4} dB, and every Monte Carlo estimate lies within 2.6 standard errors of theory. An end-to-end walk-through of a Wi-Fi link in an office, together with office and urban coverage studies, reports path-loss exponents, delay spreads, emergent Rayleigh statistics, and the effect of diffraction, and compares them with ITU-R recommendations and published measurements.**

***Index Terms*—Ray tracing, radio propagation, image method, beam tracing, ITU-R P.2040, knife-edge diffraction, fading channels, OFDM, MIMO, engineering education.**

I. Introduction

Planning a wireless link in a real building requires answers at two levels. At the propagation level, one needs to know how much power reaches each location, along which paths, and with what delays. At the link level, one needs to know what that channel does to a modulated signal and how diversity, beamforming, multiple-input multiple-output (MIMO) processing, and orthogonal frequency-division multiplexing (OFDM) respond to it. Standard texts treat these as separate topics [1]–[3], and so do most tools. Site-specific ray tracers such as Remcom's Wireless InSite [4] and NVIDIA's Sionna RT [5] produce path lists and coverage maps, while link-level simulators usually start from a statistical channel whose origin is not shown.

This work connects the two levels in a single tool. Wireless Channel Lab follows one link, by default a ceiling-mounted Wi-Fi access point serving a laptop in an office, through seven questions: where the power goes, why it fluctuates, what the multipath does across the band, how many bits each symbol can carry, and how diversity, beamforming, and MIMO counter the channel's impairments. The ray-traced channel is the input to every later step. The tool runs in any browser without installation, a full coverage map updates in about one second, and the physics core is small enough to audit line by line.

The contributions are as follows.

1) A 2.5D deterministic propagation engine with image-method reflections up to fourth order, ITU-R P.2040 single-layer slab reflection and transmission [6], knife-edge diffraction at free wall ends and convex corners [7], a ground-reflected twin for every ray, directional antennas, and wideband channel metrics (Section IV).
2) An exact beam-tracing pruning of the image tree [8], [9] that removes 97% of third-order image nodes in the test scenes (Section VII-D).
3) A link-level chain in which six modules consume the ray-traced channel directly, including its tap delays and departure angles (Sections III and V).
4) A verification suite that compares the JavaScript engine with independent NumPy/SciPy references for every mechanism, and reports Monte Carlo agreement in binomial standard errors (Section VI).
5) An end-to-end walk-through of the default link and two coverage case studies, with comparisons to ITU-R P.1238 and published measurements (Section VII).

II. Related Work

Geometrical-optics ray tracing became a radio planning tool in the early 1990s [10]. The surveys by Yun and Iskander [11] and by Fuschini *et al.* [12] describe the two main families of methods. Shooting-and-bouncing-ray (SBR) methods launch many rays and keep those that pass near the receiver. Image methods find exact specular paths by mirroring the source across each surface. Wireless InSite combines SBR with an exact-path correction or an image-based search and includes diffraction [4]. Sionna RT adds differentiability and GPU acceleration [5]. Diffraction in such tools is usually modeled with the geometrical and uniform theories of diffraction [13], [14]. The image method grows exponentially with reflection order, and beam tracing, introduced for rendering [8] and later applied to room acoustics [9], is a standard way to bound that growth.

Site-general models trade geometry for statistics. ITU-R P.1238 [15] and 3GPP TR 38.901 [16] specify distance-power laws and delay-spread distributions, grounded in measurement campaigns such as those of Seidel and Rappaport [17] and Saleh and Valenzuela [18]. The engine described here uses the material data of the current ITU-R P.2040 revision [6] and

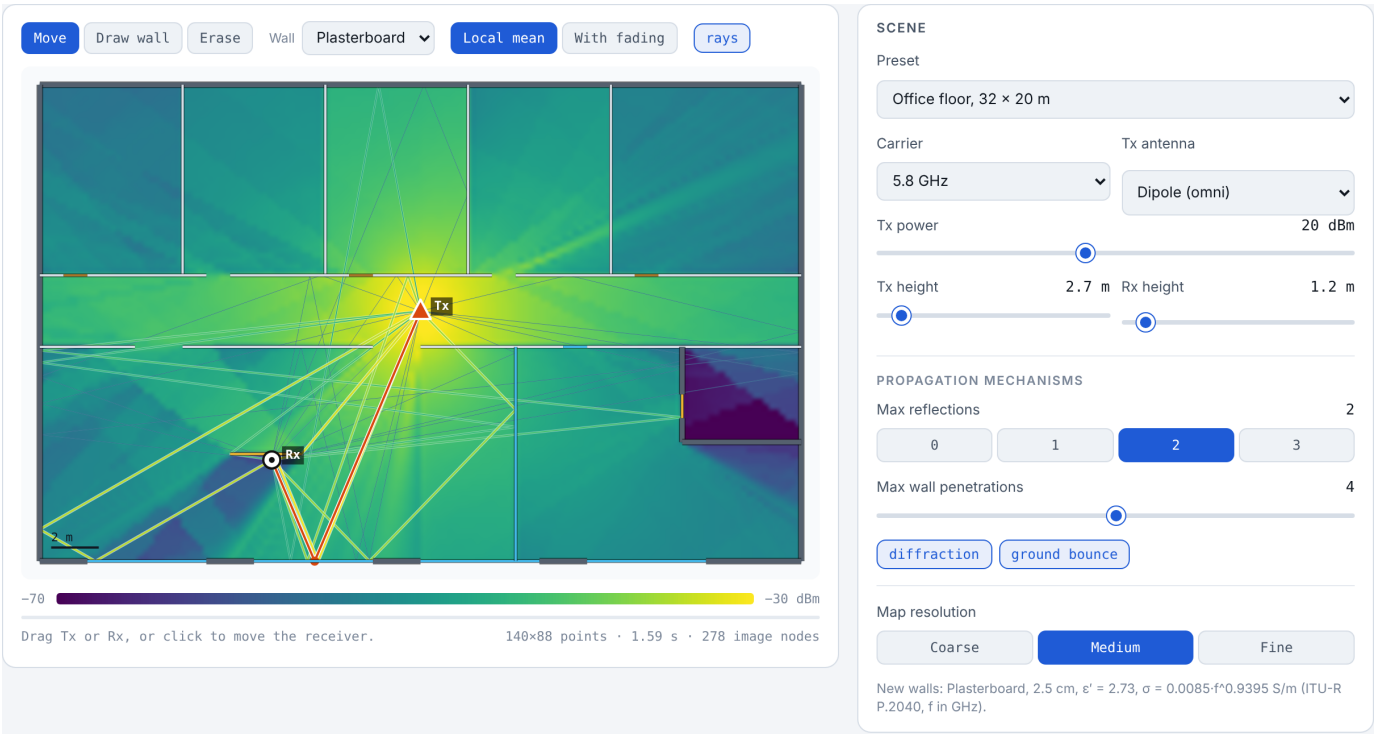


Fig. 1. The propagation module in a browser. Left: the office floor with the incoherent received-power map at 5.8 GHz, walls colored by material, and the rays reaching the laptop (Rx). The strongest ray, a single reflection from the glazed facade, is selected. Right: scene, antenna, and mechanism controls. Below the map (not shown) the module reports received power, path loss, RMS delay spread, Rician K, SNR, the power-delay profile, the angles of arrival, and a full ray table.

reports its outputs in the same terms as these models, so the two can be compared directly.

III. System Overview

Fig. 2 shows how data moves through the tool. The user edits a scene, which consists of walls with materials, antenna positions and patterns, a carrier frequency, a bandwidth, and a receiver noise figure. The propagation engine returns the full list of rays at the receiver. From that list the tool derives a compact *link record*: the local-mean received power and SNR, the Rician K-factor, the RMS delay spread, the ray taps with their delays and complex amplitudes, and the departure angle of the strongest ray. Each link-level module starts from the fields it needs and states which values it inherited. The user may override any of them, and the module then indicates that it is in manual mode until it is relinked.

IV. Propagation Engine

A. Scene geometry

A scene is a plan view of N_w wall segments, each with a material and a thickness, standing on a flat ground. The transmitter and receiver are at heights h_t and h_r, and walls are treated as infinitely tall vertical slabs. Under this geometry the three-dimensional image problem separates. The horizontal

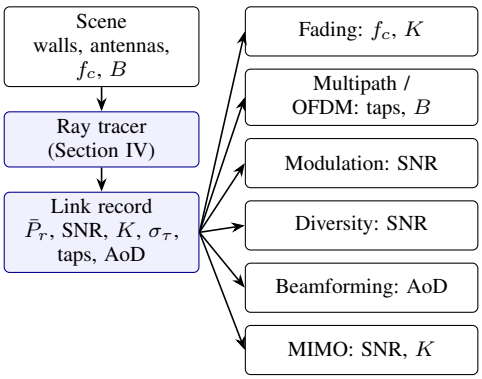


Fig. 2. Data flow in Wireless Channel Lab. The ray tracer's output is reduced to a link record, and each link-level module reads the fields it needs. The same physics core also runs under Node.js for the verification suite.

track of each ray is found in two dimensions, and its length in three dimensions is

$$L_i = \sqrt{\ell_i^2 + (h_t - h_r)^2}, \tag{1}$$

where ℓ_i is the unfolded plan-view length. Because a reflection in the ground plane commutes with reflections in vertical walls, every ray also has a ground-reflected twin of length $L_i^g = \sqrt{\ell_i^2 + (h_t + h_r)^2}$, which meets the ground at $\cos\theta_g = (h_t + h_r)/L_i^g$ from the normal. This construction is

Algorithm 1 Image tree with beam-tracing pruning

$\mathcal{F} \leftarrow \{\text{root} : \text{image} = \mathbf{x}_t,\ \text{window} = \varnothing\}$
for $k = 1$ **to** N_R **do**
  $\mathcal{F}' \leftarrow \varnothing$
  for each node $u \in \mathcal{F}$ and each wall $w \neq u.\text{wall}$ **do**
    if $k = 1$ **then** $W \leftarrow w$
    else $W \leftarrow \textsc{Clip}(w,\ u.\text{image},\ u.\text{window})$
    end if
    if $W \neq \varnothing$ **then**
      add child $(w,\ \text{mirror}(u.\text{image}, w),\ W)$ to $\mathcal{F}'$
    end if
  end for
  $\mathcal{F} \leftarrow \mathcal{F}'$
end for

exact for the stated geometry. Ceilings and finite wall heights are outside the model (Section VIII).

B. Image tree and beam-tracing pruning

The image of a point $\mathbf{p}$ across a wall through $\mathbf{w}_1$ with unit normal $\mathbf{n}$ is $\mathbf{p}' = \mathbf{p} - 2\big((\mathbf{p} - \mathbf{w}_1) \cdot \mathbf{n}\big)\mathbf{n}$. Applying this recursively to the transmitter produces an image tree whose depth-k level has up to $N_w(N_w - 1)^{k-1}$ nodes. The tree does not depend on the receiver, so it is built once for each transmitter position.

Most images cannot produce a valid path. Each node therefore stores a *window*, the part of its wall that its parent's beam can illuminate. The parent beam is the region beyond the parent's window line and between the two rays from the parent image through the window's endpoints. A child wall is kept only if part of it lies inside this beam, and the clipped segment becomes the child's window (Algorithm 1). A pruned image could not have produced a valid reflection point, so pruning leaves every result unchanged. Section VII-D confirms this numerically.

For a receiver at $\mathbf{x}_r$, each node is resolved by walking back toward the root. At each step, the segment from the node's image to the current target must cross the node's wall strictly within the segment. The crossing point becomes a reflection point and the next target. Each leg of the resulting polyline is then tested for wall crossings, which become transmissions.

C. Materials, reflection, and transmission

Wall materials follow ITU-R P.2040-3, Table 3 [6]. With f in gigahertz, the relative permittivity and conductivity are $\varepsilon_r' = a f^{\,b}$ and $\sigma = c f^{\,d}$ S/m, and the complex relative permittivity is $\eta = \varepsilon_r' - j\sigma/(2\pi f \varepsilon_0)$ with f in hertz. For an air-to-material interface at incidence angle θ, and with $s = \sqrt{\eta - \sin^2\theta}$, the Fresnel coefficients are

$$r_{\mathrm{TE}} = \frac{\cos\theta - s}{\cos\theta + s}, \qquad r_{\mathrm{TM}} = \frac{\eta\cos\theta - s}{\eta\cos\theta + s}. \tag{2}$$

A wall of thickness d is modeled as a single slab whose coefficients include all internal reflections [6]:

$$R = \frac{r\,(1 - e^{-j2q})}{1 - r^2 e^{-j2q}}, \quad T = \frac{(1 - r^2)\,e^{-j(q - q_0)}}{1 - r^2 e^{-j2q}}, \tag{3}$$

TABLE I
NORMAL-INCIDENCE TRANSMISSION AND REFLECTION LOSS (DB) OF THE DEFAULT WALLS, FROM (3) WITH ITU-R P.2040-3 PARAMETERS

	$-20\log_{10}\lvert T\rvert$ / $-20\log_{10}\lvert R\rvert$		
Material (thickness)	2.4 GHz	5.8 GHz	28 GHz
Concrete (20 cm)	14.6 / 7.9	27.5 / 8.1	90.9 / 8.1
Brick (12 cm)	3.5 / 9.7	4.0 / 10.3	5.2 / 8.1[†]
Plasterboard (2.5 cm)	1.2 / 8.0	1.9 / 7.9	5.4 / 11.2
Wood (4.5 cm)	0.7 / 30.7	2.0 / 11.0	9.0 / 16.0
Glass (1 cm)	3.1 / 3.0	0.4 / 20.6	4.2 / 5.1

[†]Outside the 1–10 GHz validity range given for brick; extrapolated.

where $q = (2\pi d/\lambda)\,s$ and $q_0 = (2\pi d/\lambda)\cos\theta$. The factor e^{jq_0} references T to the unobstructed air path, so a transmitted ray keeps its straight-line length. Antennas are vertically polarized, so the electric field is perpendicular to the horizontal plane of incidence at a vertical wall (TE) and lies in the vertical plane of incidence at the ground (TM). The ground is a half-space and uses r_{TM} directly. Table I lists the resulting normal-incidence losses for the default wall thicknesses.

The wide variation for glass and wood follows from the thin-slab resonance in (3). When the electrical thickness approaches half a wavelength inside the material, the reflections from the two faces cancel. A single loss figure per material therefore misrepresents thin walls.

Evaluating (3) at every interaction would dominate the run time. The engine instead tabulates R and T for each combination of material, thickness, and frequency at 1025 points uniformly spaced in $\cos\theta$, and interpolates linearly. Section VI shows that the resulting error is below 1.1×10^{-4} dB.

D. Diffraction

A wall endpoint is treated as a diffracting edge if all walls attached to it lie within a sector narrower than 180°; this selects free wall ends and convex corners and excludes straight joints and T-junctions. An edge contributes only if both terminals lie in its open sector, so the concave corners of a room do not leak energy outward.

Each edge is modeled as a knife edge [7]. Let d be the direct distance, d_1 and d_2 the distances from the terminals to the edge, and $\Delta = d_1 + d_2 - d$ the excess path length. With the Fresnel-Kirchhoff parameter $|\nu| = 2\sqrt{\Delta/\lambda}$, the field relative to free space is

$$F(\nu) = \frac{1 + j}{2} \int_{\nu}^{\infty} e^{-j\pi t^2/2}\, dt, \tag{4}$$

evaluated from the Fresnel integrals by a power series for $|\nu| \le 1.5$ and a continued fraction otherwise [19]. When a wall attached to the edge blocks the direct path ($\nu > 0$), the engine adds a diffracted ray $a_D = F(\nu)\,\frac{\lambda}{4\pi d}\,e^{-jkd}\prod T$, where the product contains the transmission coefficients of the two legs. When the direct path is clear but has an excess length below 6.25λ ($-5 < \nu < 0$), the direct ray is multiplied by $F(\nu)$. Beyond that range $|F(\nu)|$ is within 0.4 dB of unity. This treatment keeps the field continuous across the shadow boundary of the direct ray.

E. Field summation and link metrics

Ray i has the complex amplitude

$$a_i = \sqrt{G_t(\phi_i^d)\, G_r(\phi_i^a)}\; \frac{\lambda}{4\pi L_i}\; \Gamma_i\; e^{-jkL_i}, \tag{5}$$

where Γ_i is the product of its reflection, transmission, ground, and diffraction coefficients, and ϕ_i^d and ϕ_i^a are its departure and arrival azimuths. The available antennas are isotropic, a half-wave dipole (2.15 dBi, omnidirectional in azimuth), and the 3GPP 65° sector pattern $G(\phi) = 8 - \min\{12(\phi/65^\circ)^2, 30\}$ dBi [16]. The engine reports both the coherent and the incoherent received power,

$$P_{\mathrm{coh}} = P_t \Big| \sum_i a_i \Big|^2, \qquad P_{\mathrm{inc}} = P_t \sum_i |a_i|^2. \tag{6}$$

P_{coh} includes small-scale fading and changes over centimeters, whereas P_{inc} approximates the local mean. From the ray list the engine computes the mean excess delay $\bar{\tau}$, the power-weighted RMS delay spread $\sigma_\tau = (\overline{\tau^2} - \bar{\tau}^2)^{1/2}$ [1], the Rician K-factor, and the SNR relative to a noise floor of $-174 + 10\log_{10} B + \mathrm{NF}$ dBm.

F. Implementation

The physics core is approximately 900 lines of dependency-free JavaScript and is shared without modification by the browser and by the Node.js verification runner. Coverage maps are computed in a Web Worker that streams rows back to the page, first on a coarse grid and then at full resolution. Three modules also offer interactive 3D views rendered with three.js: the scene with extruded walls, the coverage map on the floor, and each ray drawn with its true heights, including ground bounces; the full radiation pattern of the array, whose cones around the array axis make steering and grating lobes visible; and the ergodic MIMO capacity as a surface over SNR and K. All random draws use a seeded generator, so every figure in this paper can be regenerated exactly.

V. Link-Level Modules

The modules below are presented in the order a user meets them. Each one names the fields of the link record it inherits.

Fading (inherits f_c and K). Clarke's model [20], [21] is generated as a sum of $N = 32$ sinusoids with random arrival angles and phases [22]. Random angles avoid the non-ergodic behavior of deterministic sum-of-sinusoids generators [23]. A line-of-sight term produces Rician fading [24], and the Doppler frequency is $f_D = v f_c / c$.

Multipath and OFDM (inherits the ray taps and B). The frequency response is $H(f) = \sum_\ell h_\ell e^{-j2\pi f \tau_\ell}$, evaluated either from the ray taps of the current link or from an exponential power-delay profile rescaled to a requested RMS delay spread under the WSSUS assumption [25]. The module marks subcarriers that fall more than 10 dB below the mean [26], [27] and reports the coherence-bandwidth rule $B_c \approx 1/(5\sigma_\tau)$ [1] and the required cyclic prefix.

Modulation (inherits the SNR). Symbols from BPSK, QPSK, 16-QAM, and 64-QAM constellations with unit average energy pass through complex additive white Gaussian noise (AWGN) and are detected by minimum distance. The running Monte Carlo estimate is displayed next to the exact square-QAM symbol error rate [28],

$$P_s = 1 - \Big[1 - 2\big(1 - \tfrac{1}{\sqrt{M}}\big) Q\big(\sqrt{3E_s/N_0/(M-1)}\big)\Big]^2. \tag{7}$$

Diversity (inherits the SNR). For BPSK over L independent Rayleigh branches with maximal-ratio combining (MRC) [29],

$$P_b = \Big(\tfrac{1-\mu}{2}\Big)^L \sum_{k=0}^{L-1} \binom{L-1+k}{k} \Big(\tfrac{1+\mu}{2}\Big)^k, \quad \mu = \sqrt{\tfrac{\bar{\gamma}}{1+\bar{\gamma}}} \tag{8}$$

[28], [30]. The 2×1 Alamouti code [31] performs as two-branch MRC at $\bar{\gamma}/2$.

Beamforming (inherits the departure angle). A uniform linear array with weights w_n has the array factor $AF(\theta) = \sum_n w_n e^{jn2\pi(d/\lambda)(\sin\theta - \sin\theta_0)}$ [32], [33]. The array is taken to face the receiver, and θ_0 is the angle between that bearing and the departure direction of the strongest ray. Uniform, Hamming, Blackman [34], and Dolph-Chebyshev [35] weights are provided, together with the grating-lobe condition $d/\lambda < 1/(1 + |\sin\theta_0|)$.

The module also models how the steering is implemented in hardware [36], [37]. A true-time-delay (TTD) feed delays element n by $\tau_n = x_n \sin\theta_0 / c$, which multiplies it by $e^{-j2\pi f \tau_n}$ at every frequency f, so the beam stays at θ_0 across the band. Delay lines, Rotman lenses [38], and optical delays realize this. A phase-shifter feed applies the fixed phase $\phi_n = 2\pi f_0 \tau_n$ wrapped to $[0, 2\pi)$, which is exact only at the center frequency f_0; at frequency f the beam moves to

$$\sin\theta(f) = \frac{f_0}{f}\, \sin\theta_0, \tag{9}$$

an effect known as beam squint. A B-bit phase shifter further rounds each ϕ_n to a multiple of $2\pi/2^B$. If the rounding errors are treated as uniform over $\pm\pi/2^B$, the expected gain loss is $-20\log_{10}\mathrm{sinc}(\pi/2^B)$ dB [39]. An animated view sums the cylindrical waves radiated by each element, $E(\mathbf{x}, t) = \sum_n s_n(t - r_n/c)/\sqrt{r_n}$, so the user can watch a wideband pulse leave the array on target with TTD and spread with phase shifters.

MIMO (inherits the SNR and K). The channel is $\mathbf{H} = \sqrt{K/(K+1)}\, \mathbf{a}_r \mathbf{a}_t^H + \sqrt{1/(K+1)}\, \mathbf{H}_w$, where $\mathbf{H}_w$ has independent $\mathcal{CN}(0,1)$ entries and $\mathbf{a}_r$, $\mathbf{a}_t$ are half-wavelength steering vectors. With equal power allocation the capacity is $C = \sum_k \log_2(1 + \rho\sigma_k^2/M)$ [40], [41]. The eigenvalues of $\mathbf{H}\mathbf{H}^H$ are obtained with a Jacobi solver applied to the real $2N \times 2N$ embedding of the Hermitian matrix.

VI. Verification

Each check runs the browser's JavaScript under Node.js and compares the output with a reference computed in Python from the published formula. The reference code was written from the equations rather than translated from the engine, so an error would have to be made twice in the same way to go undetected. Table II summarizes the results.

TABLE II

VERIFICATION SUMMARY. DETERMINISTIC CHECKS REPORT THE MAXIMUM ABSOLUTE ERROR. MONTE CARLO CHECKS REPORT THE LARGEST DEVIATION IN BINOMIAL STANDARD ERRORS ($|z|$) OR IN RELATIVE TERMS.

Mechanism	Reference	Result
Free space	Friis [42]	2.8×10^{-14} dB
Ground reflection	Two-ray, exact r_{TM}	1.1×10^{-4} dB
Knife edge, $\nu > -5$	Exact $F(\nu)$, SciPy	5.1×10^{-10} dB
Wall reflections, $N_R \le 4$	Analytic image sum	1.1×10^{-10} dB
Slab R and T (30 cases)	NumPy, (3)	1.1×10^{-12}
Lossless slab	$\lvert R\rvert^2 + \lvert T\rvert^2 = 1$	1.0×10^{-15}
Beam-tracing pruning	Brute-force tree	identical
Rayleigh / Rice envelope	pdf, KS distance	0.0043 / 0.0017
Level-crossing rate	$\sqrt{2\pi} f_D \rho e^{-\rho^2}$	1.5% mean
Autocorrelation	$J_0(2\pi f_D \tau)$	0.026 max.
SER, four schemes	Exact P_s [28]	$\lvert z\rvert \le 1.85$
BER, MRC and Alamouti	Closed form	$\lvert z\rvert \le 2.59$
Capacity, up to 4×4	Telatar integral [40]	1.6% max.
Squint, phase shifters	Eq. (9)	1.0×10^{-4} deg
Pointing, true time delay	$\theta = \theta_0$ for all f	1.0×10^{-4} deg

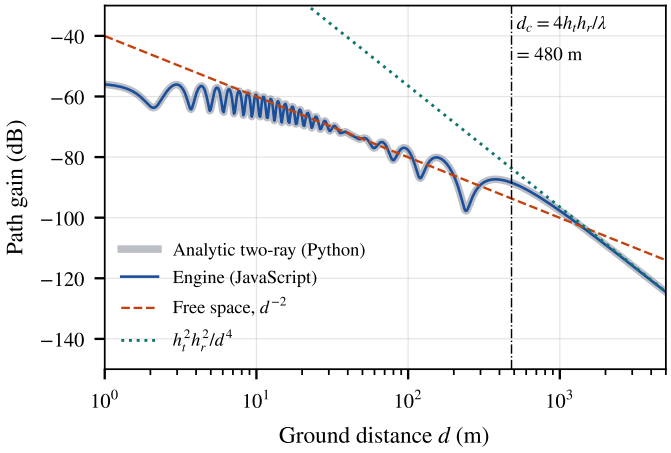


Fig. 3. Two-ray ground reflection. The engine output is drawn over the analytic two-ray field computed with the exact TM Fresnel coefficient of medium-dry ground. The d^{-4} asymptote applies beyond the break point d_c.

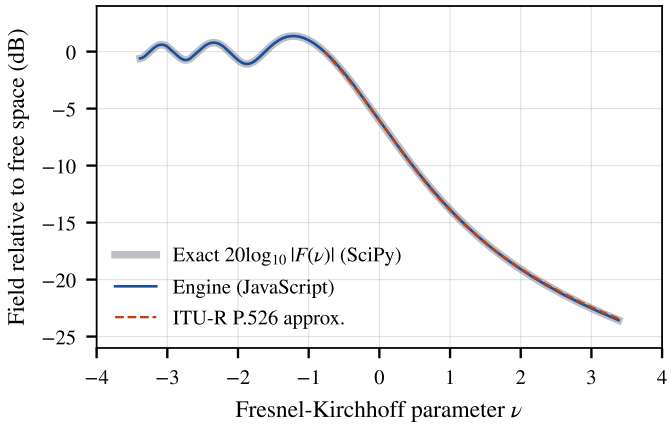


Fig. 4. Field behind a knife edge as a function of the Fresnel-Kirchhoff parameter. The engine output coincides with the exact curve. The dashed line is the ITU-R P.526 approximation, valid for $\nu > -0.78$.

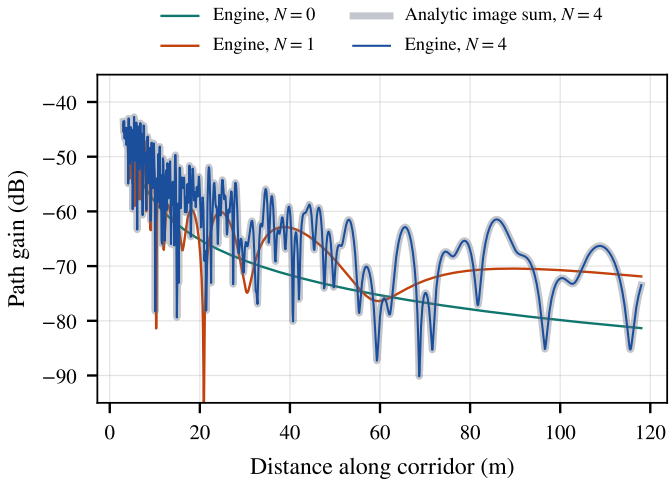


Fig. 5. Path gain along a metal corridor for maximum reflection orders $N_R = 0$, 1, and 4, with the analytic image sum for $N_R = 4$ drawn underneath.

A. Deterministic propagation mechanisms

Two-ray ground reflection. With $h_t = 10$ m, $h_r = 1.5$ m, and medium-dry ground at 2.4 GHz, the engine reproduces the analytic two-ray field to within 1.1×10^{-4} dB out to 5 km (Fig. 3). The residual is due to interpolation of the coefficient table. The curve shows the interference region, the break point near $d_c = 4h_t h_r/\lambda = 480$ m, and the d^{-4} decay beyond it.

Knife-edge diffraction. A metal half-plane is placed midway between terminals 200 m apart, with reflection and transmission disabled. Moving the receiver across the shadow boundary sweeps ν from -4 to 4 (Fig. 4). The error relative to SciPy's Fresnel integrals is 5×10^{-10} dB for $\nu > -5$, including the oscillations on the lit side and the -6.02 dB value at the shadow boundary. The closed-form approximation of ITU-R P.526 is shown for reference.

Multiple reflections. Two parallel metal walls 4 m apart form a lossy parallel-plate guide whose field is a sum over two image sequences, each with n bounces at a common incidence angle. The engine agrees with this sum to 10^{-10} dB for every maximum order from 0 to 4 (Fig. 5), which confirms that the image tree enumerates exactly the correct paths and applies the correct coefficient at each bounce.

Slab coefficients. For five materials, three frequencies, and both polarizations, the JavaScript and NumPy implementations of (3) agree to 10^{-12} in complex value, and a lossless slab satisfies $|R|^2 + |T|^2 = 1$ to machine precision. Fig. 6 shows the coefficients at 5.8 GHz.

B. Statistical modules

The fading generator was run for 40 realizations of 10 s each at $f_D = 100$ Hz. The Kolmogorov-Smirnov distance between its envelope and the target distribution is 0.0043 for Rayleigh

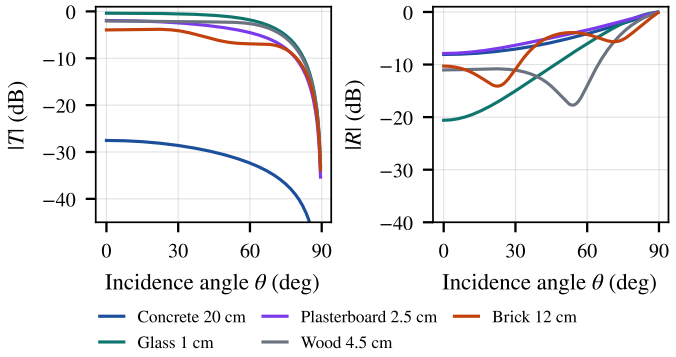


Fig. 6. TE transmission and reflection coefficients at 5.8 GHz as functions of the incidence angle for the default walls (ITU-R P.2040-3 parameters).

fading and 0.0017 for Rician fading with $K = 5$. Its level-crossing rate is within 1.5% of theory on average between -26 and $+6$ dB, and its autocorrelation follows $J_0(2\pi f_D \tau)$ (Fig. 7).

For the modulation and diversity modules, each Monte Carlo point uses enough trials to expect 300 to 400 errors, capped at 2×10^6 symbols or 3×10^6 bits. All points lie within 2.6 binomial standard errors of the closed-form result (Fig. 8), which is consistent with sampling noise over 80 points. The diversity curves show the slopes that define diversity order and the 3 dB offset between Alamouti coding and two-branch MRC.

The ergodic capacity of the MIMO module, averaged over 6000 independent Rayleigh draws, is within 1.6% of Telatar's closed-form integral [40], evaluated by quadrature with generalized Laguerre polynomials, for 1×1, 2×2, 4×2, and 4×4 links between -10 and 30 dB (Fig. 9). Increasing the line-of-sight share exposes the rank collapse that limits MIMO in sparse channels: at 20 dB the capacity of a 4×4 link falls from 22.1 b/s/Hz at $K = 0$ to 15.6 at $K = 10$ and 10.7 at $K = 100$.

For a 16-element array with half-wavelength spacing (Fig. 10), the uniform taper gives a half-power beamwidth of $6.35°$, in agreement with $0.886\lambda/(Nd)$, and a first sidelobe of -13.15 dB, which approaches the continuous-aperture value of -13.26 dB as N grows. The Dolph-Chebyshev taper holds every sidelobe at -30.00 dB with a $7.97°$ beam and a taper loss of 0.65 dB, while the Hamming taper reaches -39.4 dB with a $9.72°$ beam and a loss of 1.54 dB.

For the hardware model, the main-beam direction of a 16-element half-wavelength array was located by parabolic interpolation for steering angles of 15 to $60°$ and frequencies from $0.75 f_0$ to $1.25 f_0$ (Fig. 11). With phase shifters it follows (9) to within 10^{-4} degrees, and with true time delay it stays at θ_0 to the same precision. At a $30°$ steering angle, a phase-shifter array loses 2.4 dB toward the target at $\pm 10\%$ of f_0 and 12.6 dB at $\pm 20\%$. Averaged over steering angles from 0 to $60°$ for a 64-element array, phase quantization costs 3.80, 0.89, 0.21, and 0.05 dB for 1 to 4 bits, against 3.92, 0.91, 0.22, and 0.06 dB from the uniform-error approximation. The small shortfall arises because the rounding errors of a progressive phase are deterministic rather than uniformly random.

For the OFDM module, the exact 50%-correlation bandwidth of the exponential profile, obtained from $|R(\Delta f)| = |\sum_\ell P_\ell e^{-j2\pi \Delta f \tau_\ell}| / \sum_\ell P_\ell$ [25], lies between $0.203/\sigma_\tau$ and $0.216/\sigma_\tau$ for 8 to 64 taps. The rule $1/(5\sigma_\tau)$ is therefore accurate for this profile. For an untruncated exponential profile the exact value is $\sqrt{3}/(2\pi\sigma_\tau) \approx 0.276/\sigma_\tau$, so the rule is conservative in that case.

TABLE III
THE DEFAULT LINK TRACED THROUGH EVERY MODULE (5.8 GHz, 80 MHz, NF 7 dB)

Quantity	Value
Propagation	
Rays found / line of sight	42 / blocked
Local-mean received power	-61.8 dBm
Noise floor / average SNR	-88.0 dBm / 26.2 dB
Strongest ray	1 reflection, glazed facade
Fading and multipath	
Small-scale statistics	Rayleigh ($K = 0$)
RMS delay spread / $B_c \approx 1/(5\sigma_\tau)$	22.8 ns / 8.8 MHz
Subcarriers > 10 dB below mean	35 of 256 (13.7%)
Signal and countermeasures	
64-QAM SER in AWGN at 26.2 dB	1.6×10^{-5}
BPSK BER in Rayleigh, one antenna	6.0×10^{-4}
two-branch / four-branch MRC	$1.1{\times}10^{-6}$ / $4.6{\times}10^{-12}$
Beam steering from direct bearing	$-22°$
Ergodic capacity, $1{\times}1$ / $4{\times}4$	7.9 / 29.9 b/s/Hz

VII. Case Studies

A. End-to-end walk-through of the default link

The default scene is the office floor described in Section VII-B. A dipole access point at a height of 2.7 m in the corridor transmits 20 dBm at 5.8 GHz over an 80 MHz channel. The receiver is a laptop at a height of 1.2 m behind a free-standing metal whiteboard. Table III follows this link through all seven modules.

The table reads as a short design argument. The average SNR of 26.2 dB would support 64-QAM on a static channel, but the link has no line of sight and the rays add with unrelated phases, so the signal undergoes Rayleigh fading. A single antenna then has a bit error rate near 6×10^{-4} even for BPSK, and two-branch or four-branch combining reduces it by two or eight orders of magnitude. Across the band, a 22.8 ns delay spread gives a coherence bandwidth of about 8.8 MHz, roughly 28 Wi-Fi subcarriers, and 13.7% of the subcarriers fall more than 10 dB below the mean (Fig. 12). Those subcarriers are what coding and bit loading must protect. The strongest ray reflects from the glass facade and leaves the access point $22°$ away from the direct bearing, so a beam aimed straight at the laptop would point at a blocked path. Finally, the rich scattering that causes the fading makes the channel matrix well conditioned, and a 4×4 MIMO link at the same SNR has 3.8 times the capacity of a single antenna.

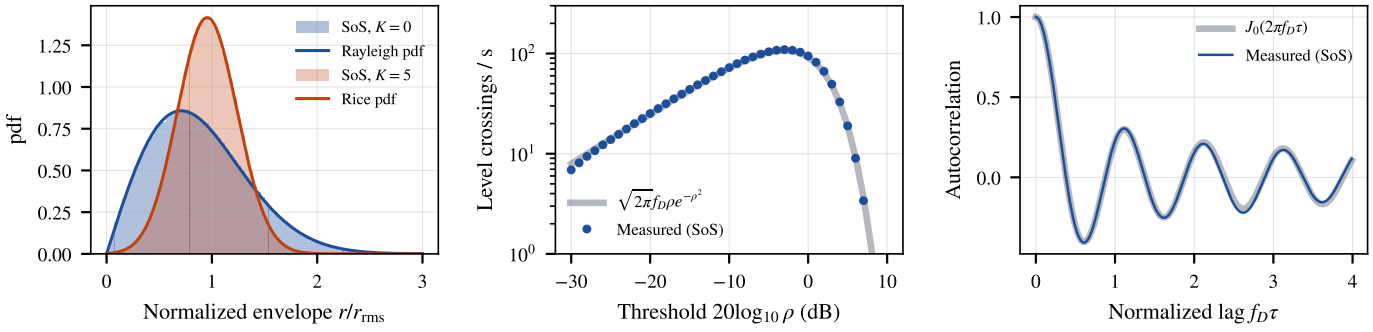


Fig. 7. Fading generator. Left: envelope histograms against the Rayleigh and Rice pdfs. Center: level-crossing rate. Right: autocorrelation.

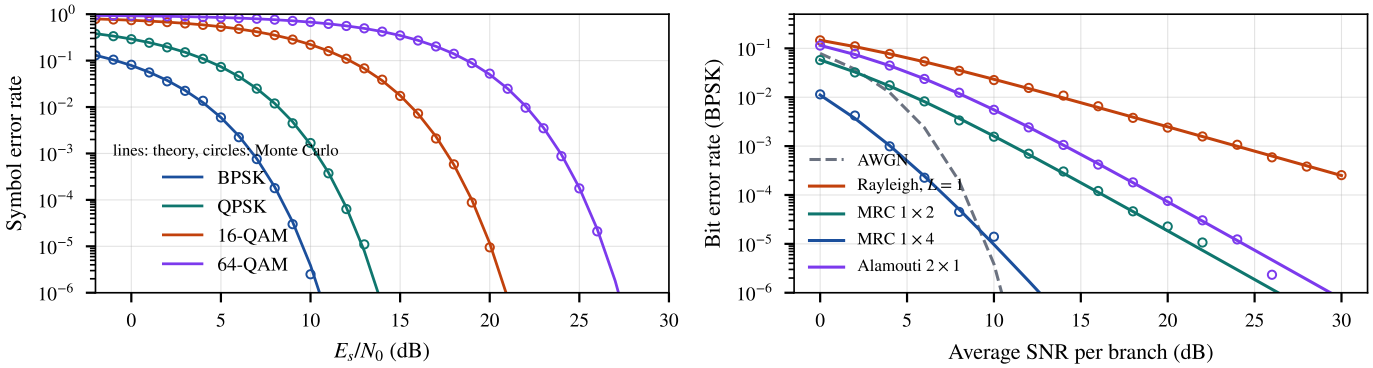


Fig. 8. Left: symbol error rate of four constellations in AWGN. Right: BPSK bit error rate with diversity in Rayleigh fading. Lines are closed-form results; circles are Monte Carlo estimates.

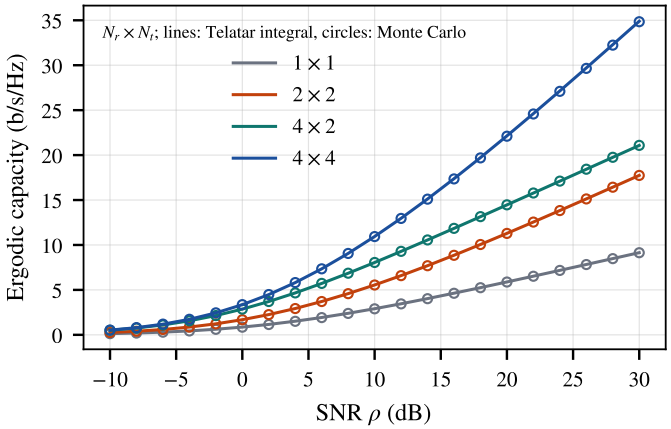


Fig. 9. Ergodic capacity of independent Rayleigh MIMO channels. Lines: Telatar's integral. Circles: Monte Carlo estimates from the browser module.

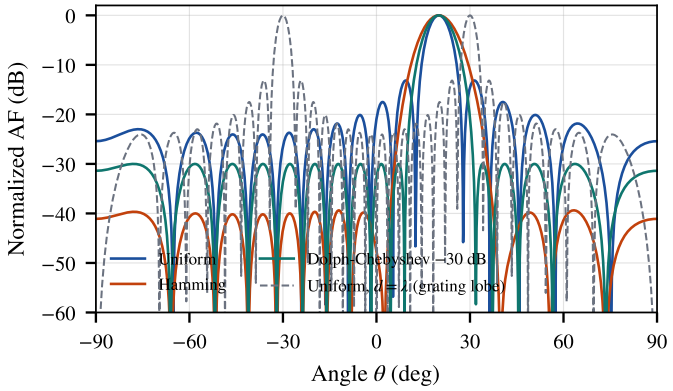


Fig. 10. Array factor of a 16-element array steered to 20° with three tapers, and the grating lobe that appears at an element spacing of one wavelength.

B. *Office floor coverage*

The office is a 32×20 m floor with a 20 cm concrete shell, a glazed south facade, plasterboard offices along a central corridor, wooden and glass doors, a concrete service core with a metal door, a free-standing metal whiteboard, and a concrete floor. To allow comparison with the 2.4 GHz literature, the coverage maps in this subsection place receivers at 1.2 m on a 12.5 cm grid (40,960 points) with $N_R = 2$ and up to four wall penetrations (Fig. 13).

The corridor behaves as a waveguide and carries the strongest signal to its ends. The offices are illuminated mainly through their door openings, and the concrete core and the metal whiteboard cast clear shadows. Moving from 2.4 to

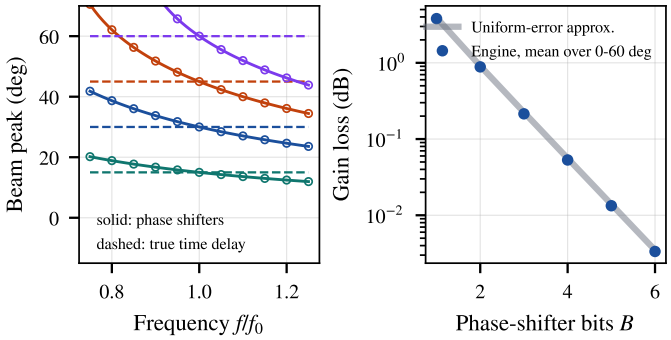


Fig. 11. Left: main-beam direction versus frequency for steering angles of 15, 30, 45, and 60°. Lines: (9); circles: engine with phase shifters; dashed: engine with true time delay. Right: gain loss from B-bit phase quantization, engine (mean over 0–60°) versus the uniform-error approximation.

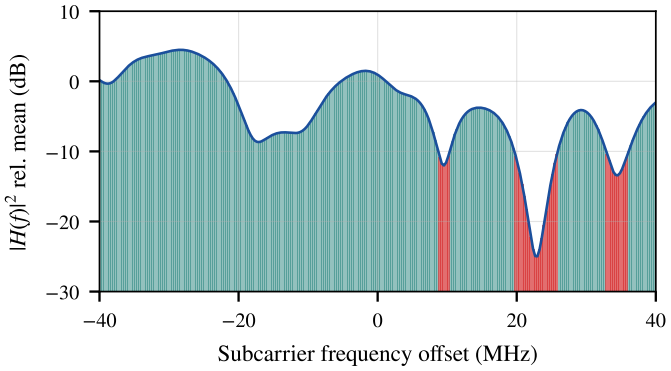


Fig. 12. Frequency response of the ray-traced default link over its 80 MHz channel, sampled at 256 subcarriers. Red bars mark subcarriers more than 10 dB below the mean.

5.8 GHz lowers the map by about 8 dB of additional free-space loss plus the higher wall losses of Table I, and the shadow of the core deepens from about −60 dBm to below −80 dBm.

A least-squares fit of $PL(d) = PL_0 + 10n\log_{10} d$ to all grid points more than 1 m from the transmitter gives $n = 2.17$ with a residual standard deviation of 4.0 dB at 2.4 GHz, and $n = 2.31$ with 4.5 dB at 5.8 GHz (Fig. 14). The fitted intercepts of 40.0 and 47.4 dB are close to the free-space loss at 1 m (40.1 and 47.7 dB). ITU-R P.1238-7 gives office distance power-loss coefficients of $N = 30$ at 2.4 GHz and $N = 31$ at 5.2 GHz [15], which correspond to $n \approx 3$. A lower exponent is expected here: the floor is small, with most points within 18 m of the access point; 24% of the floor has line of sight; the corridor guides energy; plasterboard costs only 1.2 to 1.9 dB; and the model omits furniture, people, and inter-floor propagation, all of which are present in the measurements behind P.1238.

On a 0.5 m grid at 2.4 GHz, the RMS delay spread has a median of 14.8 ns, a 90th percentile of 21.3 ns, and a maximum of 34.3 ns (Fig. 15, left). Saleh and Valenzuela reported RMS delay spreads of up to about 50 ns in a medium-size office building [18]. The simulated values fall in the same range, and the lower maximum is consistent with the smaller floor and with second-order truncation, which removes the weakest, latest arrivals.

TABLE IV
IMAGE-TREE SIZE WITH AND WITHOUT BEAM-TRACING PRUNING, AND SINGLE-THREAD RUN TIME PER RECEIVER POINT (NODE.JS 22)

	Office, $N_w = 36$				Urban, $N_w = 32$			
N_R	Full	Pruned	Ratio	µs	Full	Pruned	Ratio	µs
1	36	36	1.0	14	31	31	1.0	54
2	1293	278	4.7	77	975	190	5.1	86
3	45080	1340	33.6	285	29874	921	32.4	207
4	1.6M*	5502	290*	1183	0.98M*	3954	250*	384

*Counted analytically; the full tree at $N_R = 4$ exceeds the engine's node limit.

The right panel of Fig. 15 compares the coherent and incoherent maps point by point. In regions without line of sight the ratio $P_{\mathrm{coh}}/P_{\mathrm{inc}}$ follows the exponential distribution expected for Rayleigh fading (Kolmogorov-Smirnov distance 0.087), even though the ray tracer is fully deterministic. In line-of-sight regions the distribution is narrower, as expected for Rician fading. This observation justifies the hand-off from the propagation module to the fading module described in Section III: the statistical model assumed by the fading module emerges from the geometry once enough rays combine with unrelated phases.

C. Urban street grid

The urban scene covers 240×160 m and contains seven concrete and brick blocks, a parked bus, and medium-dry ground. A 30 dBm, 65° sector antenna at a height of 10 m points east along the main street at 3.5 GHz (Fig. 16). Without diffraction, 20.3% of the 1 m grid receives less than −120 dBm, mostly in side streets and behind blocks. With diffraction at building corners this fraction falls to 8.8%, and the side streets show the characteristic fans of diffracted energy. Inside buildings the field arrives through a single wall penetration, so values there are indicative only.

D. Convergence and run time

Table IV and Fig. 17 use fourth-order results as the reference. In the office, the mean change in local-mean power relative to fourth order is 1.51 dB at first order, 0.51 dB at second order, and 0.15 dB at third order. The urban scene converges more slowly (4.49, 1.76, and 0.39 dB) because street canyons support long multi-bounce paths. Second order is therefore a reasonable interactive default indoors, and third order is preferable outdoors.

Pruning removes 79% of the second-order tree and 97% of the third-order tree, and it changed neither the received power nor the path count at any of 600 random test points; the regression suite repeats this comparison on every change. At second order the office map costs 77 µs per point, so the 40,960-point map of Fig. 13 requires about 2.6 s on one core, and the default interactive map of about 12,000 points completes in about one second in the browser's worker thread.

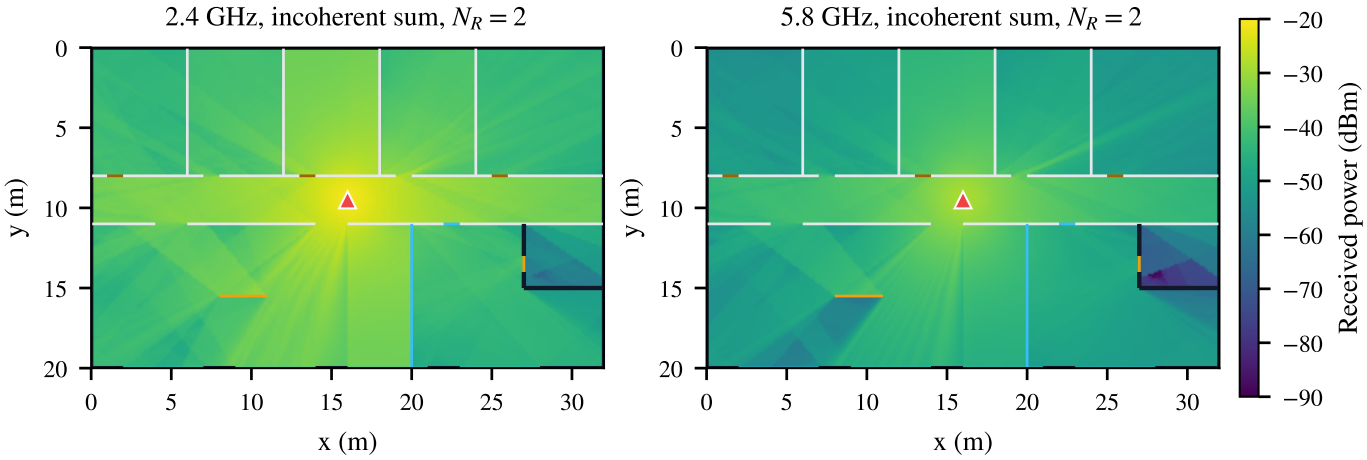


Fig. 13. Office floor, incoherent received power from a 20 dBm dipole access point (red triangle). Walls: concrete (black), plasterboard (light gray), glass (blue), wood (brown), metal (amber). Reflections up to second order, up to four penetrations, knife-edge diffraction, and floor reflection are included.

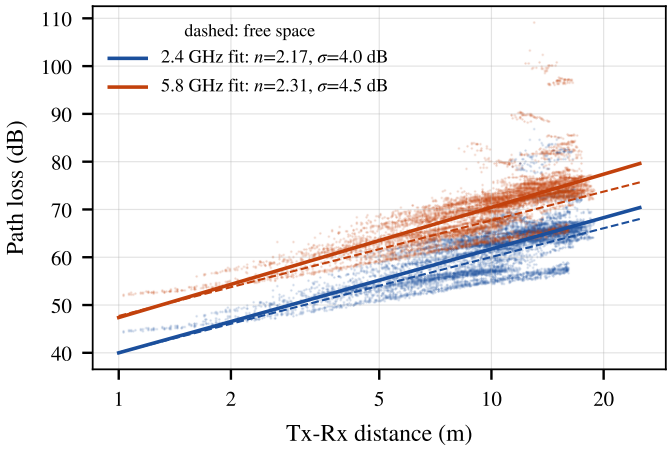


Fig. 14. Path loss versus distance over the office floor (a random 18% of the grid points shown) with least-squares fits. Dashed lines: free-space loss.

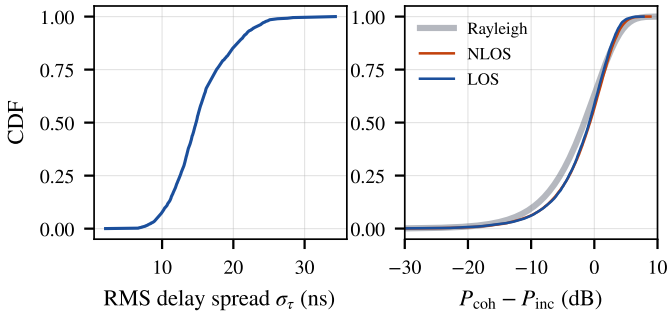


Fig. 15. Office floor at 2.4 GHz. Left: CDF of the RMS delay spread. Right: CDF of the coherent-to-incoherent power ratio in LOS and NLOS regions, compared with the Rayleigh (exponential power) CDF.

VIII. Limitations

The engine is intended for instruction and rapid what-if studies, and its simplifications are deliberate.

- *Geometry.* Walls are infinitely tall and there is no ceiling, so paths over walls and ceiling reflections are absent. Antenna patterns are isotropic in elevation, and wall incidence angles use the plan-view angle.
- *Diffraction.* The knife-edge model ignores the wedge angle and the edge material, which the uniform theory of diffraction accounts for [14], [43]. Only single diffraction of the direct path is modeled, so the field remains discontinuous at the shadow boundaries of reflected rays; these appear as straight edges in Fig. 16.
- *Mechanisms.* Diffuse scattering, furniture, and people are not modeled, and the number of wall penetrations per ray is limited.
- *Link chain.* The link-level modules use the local-mean SNR and analytical fading models; they do not simulate the time evolution of the ray-traced channel as the receiver moves.
- *Material data.* The ITU-R P.2040 fits have stated frequency ranges. Brick at 28 GHz and ground above 10 GHz are extrapolated, and the interface flags such cases.
- *Measurements.* All verification is against theory. The comparisons with P.1238 and with Saleh and Valenzuela are consistency checks, not calibrations.

Natural extensions are wedge diffraction coefficients, a ceiling and finite wall heights, a cross-check against Sionna RT on shared scenes, and a comparison with a measurement campaign.

IX. Conclusion

Wireless Channel Lab shows that a site-specific ray tracer and a set of link-level simulators can share one small, auditable core, run interactively in a browser, and still be verified

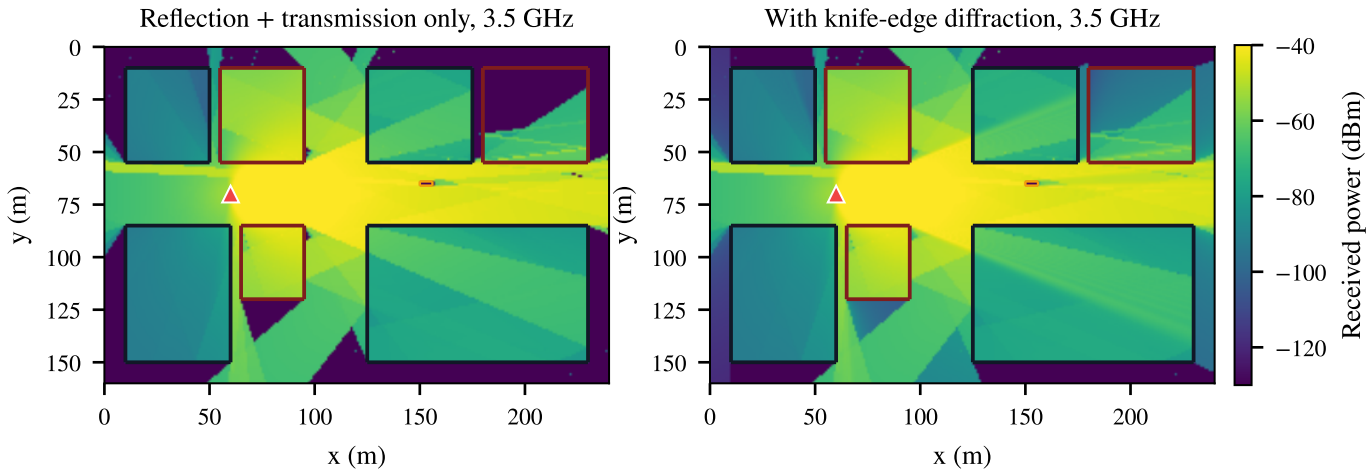


Fig. 16. Urban street grid at 3.5 GHz: incoherent received power from a sector antenna (red triangle), with and without knife-edge diffraction. Second-order reflections, one wall penetration, and ground reflection are included in both panels.

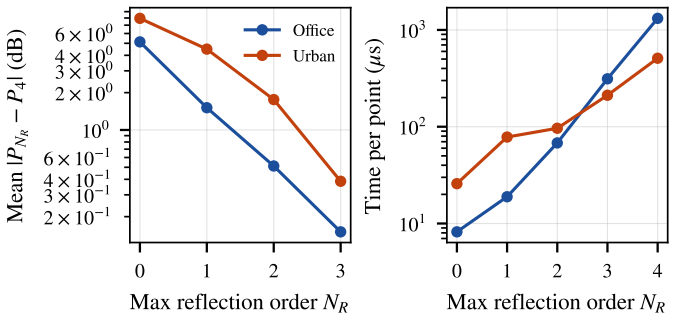


Fig. 17. Left: mean absolute change in local-mean power relative to $N_R = 4$. Right: run time per receiver point.

rigorously. Each deterministic mechanism agrees with theory to within 10^{-4} dB, and each Monte Carlo result agrees within its sampling error. Because the ray-traced channel feeds every later module, a change to the floor plan propagates to fading, error rates, beam directions, and capacity, as the walk-through of Section VII-A illustrates. The case studies reproduce established behavior, including corridor waveguiding, the two-ray break point, Rayleigh statistics emerging from deterministic rays, and diffraction filling urban shadows, and they show where a simple model departs from site-general recommendations and why. The source code, the verification suite, and the scripts that generate every figure in this paper are released under the MIT license at https://github.com/Sean-KSH/wireless-channel-lab.

References

[1] T. S. Rappaport, *Wireless Communications: Principles and Practice*, 2nd ed. Upper Saddle River, NJ, USA: Prentice Hall, 2002.

[2] A. Goldsmith, *Wireless Communications*. Cambridge, U.K.: Cambridge University Press, 2005.

[3] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge, U.K.: Cambridge University Press, 2005.

[4] Remcom, "Wireless InSite: High fidelity ray tracing," https://www.remcom.com/wireless-insite-em-propagation-software/high-fidelity-ray-tracing, accessed: Nov. 4, 2026.

[5] J. Hoydis, F. Aït Aoudia, S. Cammerer, M. Nimier-David, N. Binder, G. Marcus, and A. Keller, "Sionna RT: Differentiable ray tracing for radio propagation modeling," in *Proc. IEEE Globecom Workshops*, Kuala Lumpur, Malaysia, Dec. 2023, pp. 317–321, doi:10.1109/GCWkshps58843.2023.10465179.

[6] ITU-R, "Effects of building materials and structures on radiowave propagation above about 100 MHz," International Telecommunication Union, Geneva, Switzerland, Recommendation P.2040-3, Aug. 2023.

[7] ——, "Propagation by diffraction," International Telecommunication Union, Geneva, Switzerland, Recommendation P.526-15, Oct. 2019.

[8] P. S. Heckbert and P. Hanrahan, "Beam tracing polygonal objects," *Comput. Graph. (Proc. SIGGRAPH)*, vol. 18, no. 3, pp. 119–127, 1984, doi:10.1145/800031.808588.

[9] T. Funkhouser, I. Carlbom, G. Elko, G. Pingali, M. Sondhi, and J. West, "A beam tracing approach to acoustic modeling for interactive virtual environments," in *Proc. SIGGRAPH*, Orlando, FL, USA, Jul. 1998, pp. 21–32.

[10] J. W. McKown and R. L. Hamilton, "Ray tracing as a design tool for radio networks," *IEEE Netw.*, vol. 5, no. 6, pp. 27–30, 1991, doi:10.1109/65.103807.

[11] Z. Yun and M. F. Iskander, "Ray tracing for radio propagation modeling: Principles and applications," *IEEE Access*, vol. 3, pp. 1089–1100, 2015, doi:10.1109/ACCESS.2015.2453991.

[12] F. Fuschini, E. M. Vitucci, M. Barbiroli, G. Falciasecca, and V. Degli-Esposti, "Ray tracing propagation modeling for future small-cell and indoor applications: A review of current techniques," *Radio Sci.*, vol. 50, no. 6, pp. 469–485, 2015, doi:10.1002/2015RS005659.

[13] J. B. Keller, "Geometrical theory of diffraction," *J. Opt. Soc. Amer.*, vol. 52, no. 2, pp. 116–130, 1962, doi:10.1364/JOSA.52.000116.

[14] R. G. Kouyoumjian and P. H. Pathak, "A uniform geometrical theory of diffraction for an edge in a perfectly conducting surface," *Proc. IEEE*, vol. 62, no. 11, pp. 1448–1461, 1974, doi:10.1109/PROC.1974.9651.

[15] ITU-R, "Propagation data and prediction methods for the planning of indoor radiocommunication systems and radio local area networks in the frequency range 900 MHz to 100 GHz," International Telecommunication Union, Geneva, Switzerland, Recommendation P.1238-7, Feb. 2012.

[16] 3GPP, "Study on channel model for frequencies from 0.5 to 100 GHz," 3rd Generation Partnership Project, Technical Report TR 38.901, v17.0.0, Mar. 2022.

[17] S. Y. Seidel and T. S. Rappaport, "914 MHz path loss prediction models for indoor wireless communications in multifloored buildings," *IEEE Trans. Antennas Propag.*, vol. 40, no. 2, pp. 207–217, 1992, doi:10.1109/8.127405.

[18] A. A. M. Saleh and R. A. Valenzuela, "A statistical model for indoor multipath propagation," *IEEE J. Sel. Areas Commun.*, vol. SAC-5, no. 2, pp. 128–137, 1987, doi:10.1109/JSAC.1987.1146527.

[19] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed. Cambridge, U.K.: Cambridge University Press, 2007.

[20] R. H. Clarke, "A statistical theory of mobile-radio reception," *Bell Syst. Tech. J.*, vol. 47, no. 6, pp. 957–1000, 1968, doi:10.1002/j.1538-7305.1968.tb00069.x.

[21] W. C. Jakes, Ed., *Microwave Mobile Communications*. New York, NY, USA: Wiley, 1974.

[22] M. Pätzold, *Mobile Fading Channels: Modelling, Analysis and Simulation*. Chichester, U.K.: Wiley, 2002.

[23] M. F. Pop and N. C. Beaulieu, "Limitations of sum-of-sinusoids fading channel simulators," *IEEE Trans. Commun.*, vol. 49, no. 4, pp. 699–708, 2001, doi:10.1109/26.917776.

[24] S. O. Rice, "Mathematical analysis of random noise," *Bell Syst. Tech. J.*, vol. 23, no. 3, pp. 282–332, 1944, doi:10.1002/j.1538-7305.1944.tb00874.x.

[25] P. A. Bello, "Characterization of randomly time-variant linear channels," *IEEE Trans. Commun. Syst.*, vol. CS-11, no. 4, pp. 360–393, 1963, doi:10.1109/TCOM.1963.1088793.

[26] S. B. Weinstein and P. M. Ebert, "Data transmission by frequency-division multiplexing using the discrete Fourier transform," *IEEE Trans. Commun. Technol.*, vol. COM-19, no. 5, pp. 628–634, 1971, doi:10.1109/TCOM.1971.1090705.

[27] L. J. Cimini, Jr., "Analysis and simulation of a digital mobile channel using orthogonal frequency division multiplexing," *IEEE Trans. Commun.*, vol. COM-33, no. 7, pp. 665–675, 1985, doi:10.1109/TCOM.1985.1096357.

[28] J. G. Proakis and M. Salehi, *Digital Communications*, 5th ed. New York, NY, USA: McGraw-Hill, 2008.

[29] D. G. Brennan, "Linear diversity combining techniques," *Proc. IRE*, vol. 47, no. 6, pp. 1075–1102, 1959, doi:10.1109/JRPROC.1959.287136.

[30] M. K. Simon and M.-S. Alouini, *Digital Communication over Fading Channels*, 2nd ed. Hoboken, NJ, USA: Wiley, 2005.

[31] S. M. Alamouti, "A simple transmit diversity technique for wireless communications," *IEEE J. Sel. Areas Commun.*, vol. 16, no. 8, pp. 1451–1458, 1998, doi:10.1109/49.730453.

[32] C. A. Balanis, *Antenna Theory: Analysis and Design*, 4th ed. Hoboken, NJ, USA: Wiley, 2016.

[33] H. L. V. Trees, *Optimum Array Processing*. New York, NY, USA: Wiley, 2002.

[34] F. J. Harris, "On the use of windows for harmonic analysis with the discrete Fourier transform," *Proc. IEEE*, vol. 66, no. 1, pp. 51–83, 1978, doi:10.1109/PROC.1978.10837.

[35] C. L. Dolph, "A current distribution for broadside arrays which optimizes the relationship between beam width and side-lobe level," *Proc. IRE*, vol. 34, no. 6, pp. 335–348, 1946, doi:10.1109/JRPROC.1946.225956.

[36] R. J. Mailloux, *Phased Array Antenna Handbook*, 3rd ed. Norwood, MA, USA: Artech House, 2017.

[37] R. C. Hansen, *Phased Array Antennas*, 2nd ed. Hoboken, NJ, USA: Wiley, 2009.

[38] W. Rotman and R. F. Turner, "Wide-angle microwave lens for line source applications," *IEEE Trans. Antennas Propag.*, vol. 11, no. 6, pp. 623–632, 1963, doi:10.1109/TAP.1963.1138114.

[39] T. C. Cheston and J. Frank, "Phased array radar antennas," in *Radar Handbook*, 2nd ed., M. I. Skolnik, Ed. New York, NY, USA: McGraw-Hill, 1990, ch. 7.

[40] E. Telatar, "Capacity of multi-antenna Gaussian channels," *Eur. Trans. Telecommun.*, vol. 10, no. 6, pp. 585–595, 1999, doi:10.1002/ett.4460100604.

[41] G. J. Foschini and M. J. Gans, "On limits of wireless communications in a fading environment when using multiple antennas," *Wireless Pers. Commun.*, vol. 6, no. 3, pp. 311–335, 1998, doi:10.1023/A:1008889222784.

[42] H. T. Friis, "A note on a simple transmission formula," *Proc. IRE*, vol. 34, no. 5, pp. 254–256, 1946, doi:10.1109/JRPROC.1946.234568.

[43] R. J. Luebbers, "Finite conductivity uniform GTD versus knife edge diffraction in prediction of propagation path loss," *IEEE Trans. Antennas Propag.*, vol. AP-32, no. 1, pp. 70–76, 1984, doi:10.1109/TAP.1984.1143189.