Exposure-Constrained Beamforming with a Closed-Form Body Absorption Operator

Robin Wydaeghe*, Luc Martens, *Member, IEEE*, Günter Vermeeren, *Member, IEEE*, Emmeric Tanghe, *Member, IEEE*, and Wout Joseph, *Senior Member, IEEE*

***Abstract*—Human exposure assessment in realistic conditions requires a model of body absorption for every scene–body configuration. However, Finite-Difference Time-Domain (FDTD) simulations exceed 10^{12} cells at millimeter-wave (mmWave) frequencies. Therefore, a closed-form *exposure operator* is proposed, derived from the same ray-traced paths as the communication channel. A reusable phantom kernel separates body response from array and scene, and the method's one approximation is small, below a tenth of a percent. The operator gives the whole-body Specific Absorption Rate (wbSAR) and the surface-averaged Absorbed Power Density (APD). This quasi-exact method is over 10^6 times faster than FDTD. The numerical demonstration covers 4 anatomical phantoms, 16–256 antenna elements, eight frequencies from 8 to 100 GHz under Line-of-Sight (LOS) and Non-Line-of-Sight (NLOS) propagation, 16 scene realizations, and nine served-user positions. For a Base Station (BS) Equivalent Isotropic Radiated Power (EIRP) fixed at 55 dBm on the 256-element array under baseline Maximum-Ratio Transmission (MRT), the worst configuration reaches 4.6% of the local restriction and 3.5% of the wbSAR restriction. At half the MRT signal power, Exposure-Constrained Beamforming (ECBF) reduces median whole-body absorbed power by factors of 5.42 in LOS and 5.18 in NLOS.**

***Index Terms*—Absorbed power density, exposure-constrained beamforming, downlink MIMO, millimeter-wave, ray tracing, electromagnetic exposure compliance.**

I. Introduction

FIFTH-GENERATION (5G) and emerging sixth-generation (6G) wireless systems use hundreds of antenna elements in millimeter-wave (mmWave) arrays and steer them at nearby people [1]. Above 6 GHz, the 2020 International Commission on Non-Ionizing Radiation Protection (ICNIRP) guidelines [2] specify the basic restriction as Absorbed Power Density (APD) averaged over a 1 and/or 4 cm^2 area on the body surface, instead of Peak Spatial Specific Absorption Rate averaged over 10 g (psSAR10g). The 1 cm^2 restriction applies above 30 GHz in addition to the 4 cm^2 restriction. The whole-body Specific Absorption Rate (wbSAR) restriction also applies. A Base Station (BS) must know which applicable restriction is reached first as transmit power increases, because that restriction sets the transmit-power limit. Besides high-power downlink scenarios, realistic exposure assessment remains important as 5G and 6G combine higher antenna counts with higher frequencies [3], [4]. Moreover, accurate, rapid estimates enable exposure control during the design or operation of BSs. While commercial[1] BSs are rarely constrained by exposure limits, exposure reduction remains desirable under the As-Low-As-Reasonably-Achievable (ALARA) principle.

Hochwald and Love [5], [6] introduced the Specific Absorption Rate (SAR) matrix formulation for exposure-aware beamforming. Here, the SAR averaged over a measurement volume is written as a Hermitian quadratic form $\mathbf{x}^H\mathbf{S}\mathbf{x}$ in the precoder $\mathbf{x}$ of an M-element transmitter, with $\mathbf{S}$ the SAR matrix. In that setting the transmitter is the user's handset or User Equipment (UE). Building on [5], [6], Ying et al. [7] solve the exposure-constrained Signal-to-Noise Ratio (SNR) problem in closed form, the optimal precoder a regularizer-shifted matched filter, and Ebadi-Shahrivar et al. [8] give a compliance test whose number of measurements grows linearly with the element count. Later work extends the form to Multi-User Multiple-Input Multiple-Output (MU-MIMO) with limited Channel State Information (CSI) [9]. All of this prior work addresses uplink handset self-exposure, where the user's own device radiates against a fixed head or body geometry a few centimeters away. In the BS downlink, the body stands in the far-field of the array, and the relevant channel is the ray-traced propagation channel, not a fixed device uplink scenario.

In existing work, $\mathbf{S}$ is computed by full-wave Finite-Difference Time-Domain (FDTD) simulation on an anatomical phantom, once per antenna mounting, body posture, and array-body geometry. A mesh of ten cells per wavelength reaches order 10^8 cells at 6 GHz and 10^{12} near 100 GHz. At those frequencies the simulation becomes intractable. Castellanos et al. [10] avoid the repeated simulation with a closed-form pointwise exposure matrix at body-surface samples. Each sample is a rank-one form built from the array's free-space steering response, so it has no propagation environment. The downlink-body channel is multipath, and a single free-space direction cannot represent the phase between paths that becomes significant in the coherent regime. No closed-form, ray-

All authors are with the WAVES research group, Department of Information Technology, Ghent University–imec, Technologiepark-Zwijnaarde 126, 9052 Ghent, Belgium.

*Corresponding author. E-mail: robin.wydaeghe@ugent.be.

[1]Military systems and transmitters near occupational workers can bind on a downlink exposure limit.

traced exposure operator exists for the downlink-body problem. Such an operator must stay tractable at mmWave frequencies and physically accurate for the multipath downlink-body channel.

This paper derives the *exposure operator* $\mathbf{Q}$ in closed form from the same propagation paths that ray tracing already produces for the communication channel, with no full-wave step in the loop. $\mathbf{Q}$ is the downlink counterpart of the Hochwald-Ying SAR matrix $\mathbf{S}$. Integrating the absorbed power density over the body surface gives $\mathbf{Q}$, an $M \times M$ Hermitian positive-semidefinite matrix. The quadratic form $\mathbf{x}^H\mathbf{Q}\mathbf{x}$ is then the whole-body absorbed power under precoder $\mathbf{x}$ of length M, the number of array elements. One small approximation of the near-tissue field, below a tenth of a percent, reduces each surface point to a squared norm of an analytical Fresnel field channel. The incoherent limit is a surface absorption law of Wydaeghe et al. [11]. For the single-stream, one-active-operator problem, the closed-form Exposure-Constrained Beamforming (ECBF) [7] works unchanged with the analytical $\mathbf{Q}$ in place of the FDTD-computed matrix.

To the best of the authors' knowledge, the following contributions are novel.

1) The first closed-form exposure operator for downlink body exposure, enabling rapid assessment and control of exposure. It is assembled from the same ray-traced paths as the communication channel, with the body reduced to a phantom-only absorption kernel and no FDTD simulations required. The kernel is precomputed once per phantom and reused across scenes and arrays.
2) The operator is exact and positive semidefinite, and remains so under refinements of the surface response such as multilayer skin, shadow-edge diffraction, and near-field illumination. A small approximation, below a tenth of a percent, reduces it to a squared norm with the diagonal exact. In the incoherent limit it reduces to a full-wave-validated surface law.
3) A formulation of ECBF for the general downlink problem with every applicable ICNIRP 2020 basic restriction active at once. The closed-form ECBF precoder [7] works unchanged with the analytical $\mathbf{Q}$ in place of the FDTD-computed matrix.
4) An indoor-factory demonstration of the operator across the Frequency Range 2 (FR2) and Frequency Range 3 (FR3) bands, covering 27 648 scene–body–array items and 746 496 full-body maps.

The remainder of this paper is organized as follows. Section II fixes notation and assembles the field, signal, and exposure quantities from one ray-traced channel. Section III derives the coherent absorption law. Section IV reduces the body to the absorption kernel and factors the operator. Section V discusses the ECBF problem. Section VI reports the ray-traced indoor-factory demonstration across the FR2 and FR3 bands, quantifies the exposure reduction relative to Maximum-Ratio Transmission (MRT), and identifies the binding restriction for each body. Finally, Section VII concludes.

II. System model

Fig. 1 sets the scene: a ray-traced indoor factory used for the numerical demonstration in Section VI. A mast-mounted Uniform Rectangular Array (URA) at 28 GHz serves a device beside a body phantom among the factory scatterers. The two propagation scenarios use the same scene: the direct path is open under Line-of-Sight (LOS), whereas the red metallic blocker closes it under Non-Line-of-Sight (NLOS). The insets preview the field distributions produced by MRT and ECBF.

The ray-traced path model uses the following notation. Boldface lowercase and uppercase letters denote column vectors and matrices, $(\cdot)^T$ and $(\cdot)^H$ transpose and conjugate transpose, $\mathbf{I}$ the identity, and a hat a unit vector. The array has M elements and radiates with precoder $\mathbf{x} = [x_1, \dots, x_M]^T \in \mathbb{C}^M$, normalized so that $\|\mathbf{x}\|^2 = P$ is the total transmit power. The propagation environment between the array and the body produces N paths through reflection, scattering, and diffraction. Path n, for $n = 1, \dots, N$, carries the triple $(j(n), \hat{\boldsymbol{k}}_n, \boldsymbol{\psi}_n)$. Here, $j : \{1, \dots, N\} \to \{1, \dots, M\}$ maps each path to its originating element, $\hat{\boldsymbol{k}}_n$ is its direction of arrival at the body surface, and $\boldsymbol{\psi}_n \in \mathbb{C}^3$ is its complex polarization-amplitude vector orthogonal to $\hat{\boldsymbol{k}}_n$. That vector encodes the path amplitude at the body, the polarization state on arrival, and every propagation interaction along its path. A ray tracer, e.g. Sionna RT [12], produces $(j(n), \hat{\boldsymbol{k}}_n, \boldsymbol{\psi}_n)$ as its native output.

In the downlink high-frequency scope of our method, three assumptions are warranted. First, the body surface Σ is locally flat on the scale of the wavelength, so Fresnel theory applies pointwise. Local flatness weakens on the smallest features, where the curvature radius approaches the wavelength, and the curvature refinement of Section IV improves it. Second, the penetration depth is much smaller than any body feature. At 28 GHz the power penetration depth is $\delta \sim 0.5$ mm, against curvature radii of ~ 50 mm on the torso and ~ 8 mm on the smallest features such as ears and fingers, so all power transmitted through the surface is absorbed within a thin surface layer. Third, each path reaches the body as a single plane wave: its direction $\hat{\boldsymbol{k}}_n$ and polarization $\boldsymbol{\psi}_n$ are common to every surface point, and only the geometric phase $e^{-ik_0\hat{\boldsymbol{k}}_n\cdot\mathbf{r}}$ varies across Σ. The plane-wave assumption holds when the last-scatter point lies far from the body compared with its extent, so that wavefront curvature and the spread of arrival directions across Σ are negligible. The array-to-body range in Section VI is 6 to 22 m. Reflected paths additionally require their last interaction point to remain far from the body compared with its extent. Near-field illumination lies outside the present factorization, as discussed in Section IV.

Three matrices describe the ray-traced paths. First, the element-to-path dispatch matrix is

$$\mathbf{J} \in \{0,1\}^{N\times M}, \qquad J_{nm} = \delta_{j(n),m}. \tag{1}$$

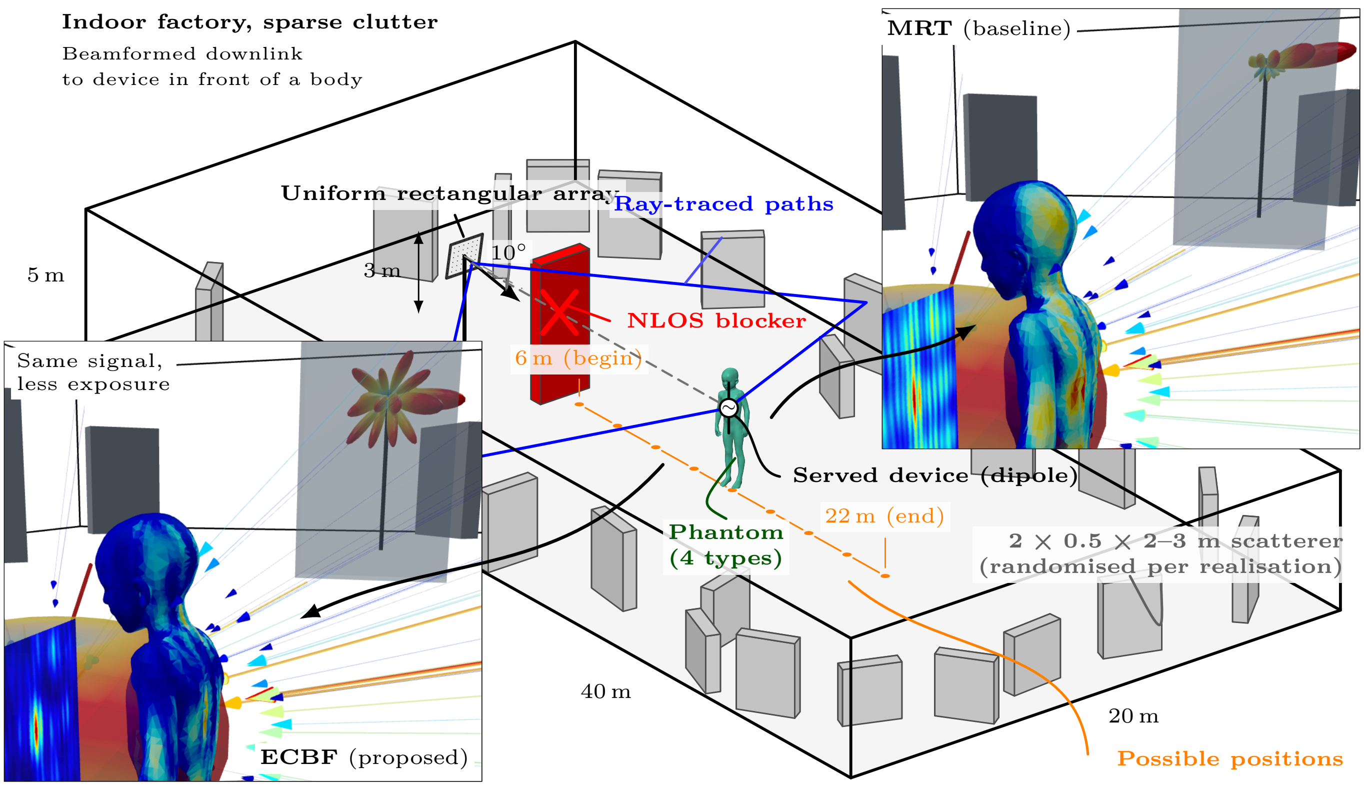

Fig. 1. The indoor-factory configuration and beamforming effect. A mast-mounted uniform rectangular array with 10° downtilt serves nine device positions over 6–22 m. A body phantom stands behind the served dipole, and a metallic blocker blocks the direct path for the NLOS condition. The insets compare MRT and ECBF at the same received signal power: MRT concentrates the transmitted field, whereas ECBF reduces body exposure.

Second, for an observation point $\mathbf{r} \in \mathbb{R}^3$, the phase diagonal matrix is

$$\boldsymbol{\Phi}(\mathbf{r}) = \operatorname{diag}\left(e^{-ik_0\hat{\boldsymbol{k}}_1\cdot\mathbf{r}}, \ldots, e^{-ik_0\hat{\boldsymbol{k}}_N\cdot\mathbf{r}}\right), \tag{2}$$

a unitary diagonal that applies the geometric phase of each path at $\mathbf{r}$, with free-space wavenumber $k_0 = 2\pi/\lambda$. Third, the polarization matrix stacks the per-path polarization-amplitude vectors as columns,

$$\boldsymbol{\Psi} = [\boldsymbol{\psi}_1, \ldots, \boldsymbol{\psi}_N] \in \mathbb{C}^{3\times N}. \tag{3}$$

$\boldsymbol{\Psi}$ depends on the propagation environment and the transmit-antenna patterns. It depends on neither the precoder nor the observation point.

The free-space electric field at any point $\mathbf{r}$ under precoder $\mathbf{x}$ is the sum of the per-path contributions,

$$\mathbf{E}(\mathbf{r}) = \sum_n \boldsymbol{\psi}_n\, x_{j(n)}\, e^{-ik_0\hat{\boldsymbol{k}}_n\cdot\mathbf{r}} = \mathbf{G}(\mathbf{r})\,\mathbf{x}, \tag{4}$$

with field channel $\mathbf{G}(\mathbf{r}) = \boldsymbol{\Psi}\,\boldsymbol{\Phi}(\mathbf{r})\,\mathbf{J} \in \mathbb{C}^{3\times M}$. In this factorization $\boldsymbol{\Psi}$ contains the environment, $\boldsymbol{\Phi}$ the position, and $\mathbf{J}$ the dispatch. Define the per-path excitation at $\mathbf{r}$,

$$\mathbf{c}(\mathbf{r}, \mathbf{x}) = \boldsymbol{\Phi}(\mathbf{r})\,\mathbf{J}\,\mathbf{x} \in \mathbb{C}^N. \tag{5}$$

Entry n of $\mathbf{c}$ is the precoder weight of the originating element of path n times the geometric phase of that path at $\mathbf{r}$. This excitation is the state variable from which every coherent observable follows. The free-space field is $\mathbf{E}(\mathbf{r}) = \boldsymbol{\Psi}\,\mathbf{c}(\mathbf{r}, \mathbf{x})$, and the UE signal is $y(\mathbf{x}) = \mathbf{a}^T\mathbf{c}(\mathbf{r}_{\mathrm{UE}}, \mathbf{x}) = \mathbf{h}^T\mathbf{x}$, with $\mathbf{h} \in \mathbb{C}^M$ the signal channel and per-path antenna projections $a_n = \mathbf{C}_R(\hat{\boldsymbol{k}}_n)^H\boldsymbol{\psi}_n$ through the UE antenna response $\mathbf{C}_R(\hat{\boldsymbol{k}}) \in \mathbb{C}^3$ in direction $\hat{\boldsymbol{k}}$. The exposure channel of Section III extends the same structure to the body surface.

III. Coherent absorption

A. Per-path transmitted field

Fig. 2 shows a ray entering the tissue of a human body Σ. Using geometric optics, we form the *exposure channel* from the same ray-traced paths as the signal channel. Each ray-traced path is a local plane wave that refracts through the surface by Fresnel theory and is absorbed in the tissue beneath. The transmitted fields are summed coherently rather than as powers, which keeps the interference between paths, and diffraction and curvature enter as refinements in Section IV. Fig. 3 shows a flowchart with the signal and exposure branches before they meet in the constrained precoder.

At a surface point $\mathbf{r}$ with outward normal $\hat{\boldsymbol{n}}$, path n arrives with incidence cosine $\mu_n = \hat{\boldsymbol{n}}\cdot(-\hat{\boldsymbol{k}}_n)$. For a front-facing path ($\mu_n > 0$) the local Transverse Electric (TE) and Transverse Magnetic (TM) unit vectors are $\hat{e}_{s,n} = \hat{\boldsymbol{k}}_n\times\hat{\boldsymbol{n}}/|\hat{\boldsymbol{k}}_n\times\hat{\boldsymbol{n}}|$ and $\hat{e}_{p,n} = \hat{e}_{s,n}\times\hat{\boldsymbol{k}}_n$. The Fresnel amplitude transmission coefficients are

$$t_{s,n} = \frac{2\mu_n}{\mu_n + \xi_n}, \quad t_{p,n} = \frac{2\tilde{n}\mu_n}{\tilde{n}^2\mu_n + \xi_n}, \tag{6}$$

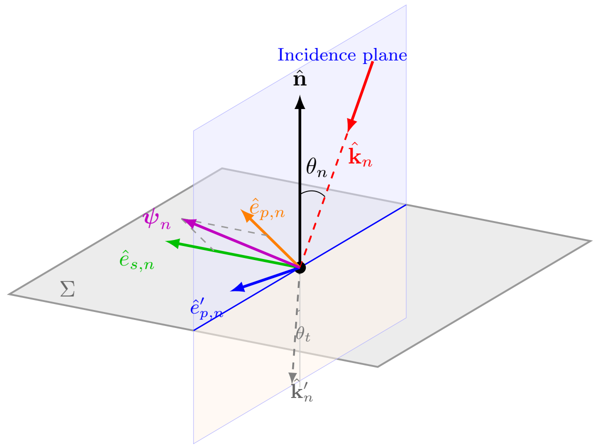

Fig. 2. Incidence-plane geometry for path n. The polarization amplitude $\boldsymbol{\psi}_n$ is resolved on the TE basis $\hat{e}_{s,n}$ (out of plane) and the incident TM projection basis $\hat{e}_{p,n}$ (in plane). The transmitted TM vector is the exact complex vector $\mathbf{e}'_{p,n} = (\xi_n \hat{\boldsymbol{t}}_n + \sin\theta_n\,\hat{\boldsymbol{n}})/\tilde{n}$, where $\hat{\boldsymbol{t}}_n$ is the in-plane tangent and ξ_n uses the passive branch. It is not unit-normalized.

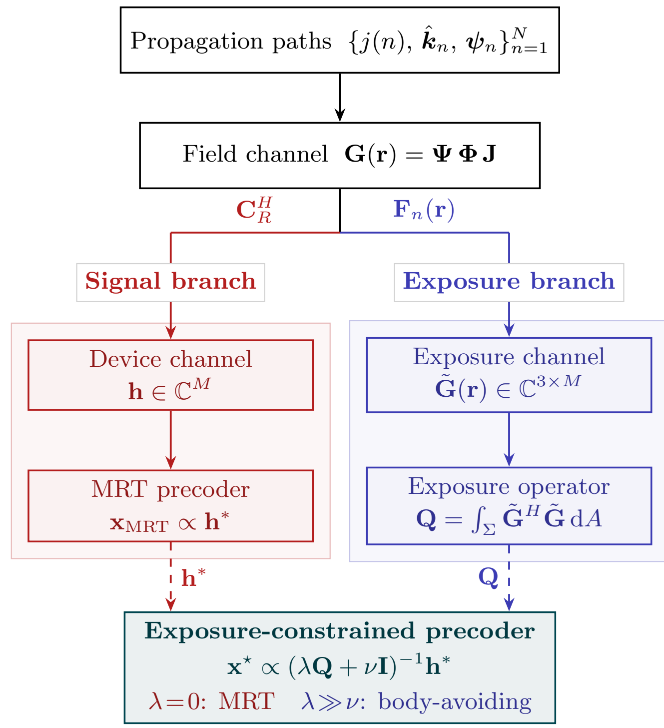

Fig. 3. Flowchart of the shared signal and exposure channel. The branches share the path dispatch and phase, but use different amplitude blocks. The signal branch gives $\mathbf{h}$ and MRT. The exposure branch gives $\tilde{\mathbf{G}}(\mathbf{r})$ and $\mathbf{Q}$. The branches meet in the constrained precoder $\mathbf{x}^\star \propto (\lambda\mathbf{Q}+\nu\mathbf{I})^{-1}\mathbf{h}^*$, with Lagrange multipliers λ (exposure) and ν (power). The precoder steers power away from the body as λ grows.

with $\xi_n = \sqrt{\tilde{n}^2 - 1 + \mu_n^2}$ and $\mathrm{Re}(\xi_n) > 0$. Writing $k_0\xi_n = \beta_n - i\alpha_n$, where $\alpha_n = -k_0\,\mathrm{Im}(\xi_n) > 0$ is the depth-attenuation constant, the depth factor inside tissue is $e^{-ik_0\xi_n z} = e^{-\alpha_n z}e^{-i\beta_n z}$. The transmitted field points along the refracted directions. The TE part keeps $\hat{e}_{s,n}$, and the TM part follows the refracted direction $\hat{e}'_{p,n}$, which the high tissue index tilts to within $O(1/|\tilde{n}|)$ of the surface tangent. The transmitted electric field of path n at depth z is then

$$\mathbf{E}_n^{\mathrm{trans}}(\mathbf{r}, z) = \big(t_{s,n}\psi_{s,n}\hat{e}_{s,n} + t_{p,n}\psi_{p,n}\hat{e}'_{p,n}\big) \times x_{j(n)}\, e^{-ik_0\hat{\boldsymbol{k}}_n\cdot\mathbf{r}}\, e^{-ik_0\xi_n z}\, V(\mathbf{r}, \hat{\boldsymbol{k}}_n)\, H(\mu_n), \tag{7}$$

with $\psi_{s,n} = \hat{e}_{s,n}\cdot\boldsymbol{\psi}_n$ and $\psi_{p,n} = \hat{e}_{p,n}\cdot\boldsymbol{\psi}_n$ the incident amplitudes, and H the Heaviside step function that sets back-facing paths to zero.

The factor $V(\mathbf{r}, \hat{\boldsymbol{k}}_n) \in \{0, 1\}$ is the binary self-shadowing visibility for path n at $\mathbf{r}$, one when an unobstructed ray reaches $\mathbf{r}$ from $\hat{\boldsymbol{k}}_n$ and zero when another body part occludes it. This is the ambient-occlusion term of the computer-graphics literature [13]–[15], and it removes self-shadowed contributions from the coherent sum. The amplitude, visibility factors, and transmitted directions define a per-path Fresnel transmission operator $\mathbf{F}_n(\mathbf{r})$ that maps the incident polarization $\boldsymbol{\psi}_n$ to the transmitted surface field $\mathbf{F}_n(\mathbf{r})\boldsymbol{\psi}_n$, rank two and set to zero on self-shadowed paths. The total field is the coherent sum $\mathbf{E}^{\mathrm{trans}} = \sum_n \mathbf{E}_n^{\mathrm{trans}}$, and the APD at $\mathbf{r}$ is the ohmic loss integrated over depth,

$$S_{\mathrm{ab}}(\mathbf{r}) = \frac{\sigma}{2}\int_0^\infty |\mathbf{E}^{\mathrm{trans}}(\mathbf{r}, z)|^2\,\mathrm{d}z. \tag{8}$$

B. Exact quadratic form

Substituting the coherent sum into Eq. (8) and carrying out the depth integral separates a depth factor, which depends only on the attenuation and phase of each path pair, from a surface factor, which depends only on the transmitted fields at $z = 0$. The depth factor is

$$\Lambda_{nn'} = \int_0^\infty e^{-(\alpha_n+\alpha_{n'})z}\, e^{-i(\beta_{n'}-\beta_n)z}\,\mathrm{d}z = \frac{1}{\alpha_n + \alpha_{n'} + i(\beta_{n'} - \beta_n)}. \tag{9}$$

The matrix Λ is Hermitian positive-semidefinite, with diagonal $\Lambda_{nn} = 1/(2\alpha_n)$ recovering the single-path absorbed power and off-diagonal entries holding the coherent cross-terms. The surface factor is the Gram matrix

$$D_{nn'}(\mathbf{r}) = \big(\mathbf{F}_n(\mathbf{r})\,\boldsymbol{\psi}_n\big)^H \big(\mathbf{F}_{n'}(\mathbf{r})\,\boldsymbol{\psi}_{n'}\big) \tag{10}$$

of the transmitted surface fields, Hermitian positive-semidefinite by construction. Writing the per-path excitation $\mathbf{c}(\mathbf{r}) = \boldsymbol{\Phi}(\mathbf{r})\,\mathbf{J}\,\mathbf{x}$ of Section II, with entries $c_n = x_{j(n)}\, e^{-ik_0\hat{\boldsymbol{k}}_n\cdot\mathbf{r}}$, the absorbed power density is the Hermitian quadratic form

$$S_{\mathrm{ab}}(\mathbf{r}) = \frac{\sigma}{2}\,\mathbf{c}(\mathbf{r})^H \big(\Lambda \circ D(\mathbf{r})\big)\,\mathbf{c}(\mathbf{r}), \tag{11}$$

where $\circ$ is the entrywise product. The entrywise product of two positive-semidefinite matrices is positive-semidefinite by the Schur product theorem, so $S_{\mathrm{ab}}(\mathbf{r}) \geq 0$ for every precoder with no approximation.

C. Squared-norm reduction

The reduction normalizes Λ by its diagonal,

$$\Gamma_{nn'} \equiv \frac{\Lambda_{nn'}}{\sqrt{\Lambda_{nn}\,\Lambda_{n'n'}}} = \frac{2\sqrt{\alpha_n \alpha_{n'}}}{\alpha_n + \alpha_{n'} + i(\beta_{n'} - \beta_n)}, \quad (12)$$

so that $\Gamma_{nn} = 1$ by construction. The reduction sets $\Gamma_{nn'} = e^{-i(\beta_{n'}-\beta_n)\,\delta}$, with $\delta = 1/(2\alpha_0)$ the normal-incidence power penetration depth of Section II. This replaces $\Lambda_{nn'}$ by an outer product, exact on the diagonal and rank one off it: path n keeps its depth-integrated amplitude $\sqrt{\Lambda_{nn}}$, with its phase referenced at the depth δ rather than at the surface. The depth δ is the centroid of the power decay profile, which is where the integrand of Eq. (9) samples the dephasing of the pair, so the phase error of the reduction cancels to first order. The residual is second order and small because the tissue index is high. At 28 GHz, skin has $|\tilde{n}| \approx 4.83$, which keeps every refraction angle within $13°$ of normal and leaves α_n and β_n nearly angle-independent, so the depth response varies negligibly from path to path. Over $(\theta_n, \theta_{n'}) \in [0°, 85°]^2$ the error $|\Gamma_{nn'} - e^{-i(\beta_{n'}-\beta_n)\delta}|$ stays below 0.05% with mean 0.014%, and below 0.08% across the band.

With the depth coupling rank one, the surface field of each path is a single tissue-weighted vector. Define

$$\tilde{\boldsymbol{\psi}}_n(\mathbf{r}) \equiv \sqrt{\frac{\sigma}{4\alpha_n}}\; e^{-i\beta_n\delta}\, \mathbf{F}_n(\mathbf{r})\, \boldsymbol{\psi}_n \in \mathbb{C}^3, \quad (13)$$

stacked as $\tilde{\boldsymbol{\Psi}}(\mathbf{r}) = [\tilde{\boldsymbol{\psi}}_1(\mathbf{r}), \dots, \tilde{\boldsymbol{\psi}}_N(\mathbf{r})] \in \mathbb{C}^{3\times N}$. The body-surface field channel then factors as

$$\tilde{\mathbf{G}}(\mathbf{r}) = \tilde{\boldsymbol{\Psi}}(\mathbf{r})\, \boldsymbol{\Phi}(\mathbf{r})\, \mathbf{J} \in \mathbb{C}^{3\times M}, \quad (14)$$

in exact parallel with the free-space field channel $\mathbf{G}(\mathbf{r}) = \boldsymbol{\Psi}\,\boldsymbol{\Phi}\,\mathbf{J}$ of Section II. Signal, field, and exposure channels share the dispatch matrix $\mathbf{J}$ and the phase matrix $\boldsymbol{\Phi}$ (Fig. 3), differing only in the per-path amplitude block: $\boldsymbol{\Psi}$ for the free-space field, $\mathbf{a}^T$ for the UE signal, and $\tilde{\boldsymbol{\Psi}}$ for the body-surface exposure. The signal branch samples one point through the UE antenna, and the exposure branch samples the whole body surface through the per-path Fresnel operator.

The APD is then the squared norm

$$S_{\text{ab}}(\mathbf{r}) = \|\tilde{\mathbf{G}}(\mathbf{r})\,\mathbf{x}\|^2 = \mathbf{x}^H\, \tilde{\mathbf{G}}(\mathbf{r})^H\, \tilde{\mathbf{G}}(\mathbf{r})\,\mathbf{x}, \quad (15)$$

with rank at most three at each surface point. The squared norm follows from the exact form Eq. (11). The reduction of Section III-C makes $\Lambda_{nn'} = \sqrt{\Lambda_{nn}}\, e^{i\beta_n\delta}\, \sqrt{\Lambda_{n'n'}}\, e^{-i\beta_{n'}\delta}$ an outer product, so the Schur product $\Lambda \circ D(\mathbf{r})$ becomes the Gram matrix of the tissue-weighted vectors of Eq. (13). The quadratic form of a Gram matrix is a squared norm, and the inner sum is $\tilde{\mathbf{G}}(\mathbf{r})\,\mathbf{x}$.

D. Limiting cases

Two limiting cases of Eq. (15) are known results. Both depend only on the exact diagonal, so they hold with or without the reduction. For a single path, Eq. (15) reduces to the polarization-aware Fresnel surface law at a visible surface point, and under the additional pseudo-Brewster approximation to the geometric law $S_{\text{ab}}(\mathbf{r}) = S_{\text{inc}}\, T_0\, V(\mathbf{r}, \hat{\boldsymbol{k}})\, [\mu]_+$, with T_0 the normal-incidence power transmission coefficient. When the N paths carry mutually random relative phases, the cross-terms average to zero and the expected APD is

$$\langle S_{\text{ab}}(\mathbf{r})\rangle = (2Z_0)^{-1} \sum_n |x_{j(n)}|^2 \big(T_{s,n}|\psi_{s,n}|^2 + T_{p,n}|\psi_{p,n}|^2\big) \times V(\mathbf{r}, \hat{\boldsymbol{k}}_n)\, [\mu_n]_+, \quad (16)$$

where $T_{s,n}$ and $T_{p,n}$ are the power transmission coefficients of the two polarizations at the incidence angle of path n. This is the standard multi-source surface dosimetry law of Wydaeghe et al. [11]. Because the self-terms are exact, this limiting law has no approximation error. Full-wave validation of the incoherent law therefore validates only the incoherent limit, not the full coherent operator. The supplementary material uses an analytical sphere as a separate consistency test for the coherent reduction.

IV. Absorption kernel

Integrating Eq. (15) over the body surface gives the whole-body absorbed power as a quadratic form in the precoder. The exposure operator is

$$\mathbf{Q} = \int_\Sigma \tilde{\mathbf{G}}(\mathbf{r})^H\, \tilde{\mathbf{G}}(\mathbf{r})\, \mathrm{d}A \in \mathbb{C}^{M\times M}, \quad (17)$$

so that $P_{\text{abs}} = \mathbf{x}^H\, \mathbf{Q}\, \mathbf{x}$ is the whole-body absorbed power under precoder $\mathbf{x}$. By construction, $\mathbf{Q}$ is Hermitian positive-semidefinite, with rank at most $\min(M, 3M_\Sigma)$, where M_Σ counts the integration triangles on the body mesh. The trace $\mathrm{tr}(\mathbf{Q})$ equals the expected absorbed power under isotropic random precoding, scaled by the transmit power per element.

$\mathbf{Q}$ is the downlink counterpart of the Hochwald-Ying SAR matrix $\mathbf{S}$ [6], [7], integrated over the body surface rather than a fixed tissue volume, so it returns whole-body absorbed power directly. An even coarser reduction keeps only the diagonal of $\mathbf{Q}$ and discards every off-diagonal entry. The discarded off-diagonal entries are the coherent coupling between paths, and Section VI shows most of the operator energy is in them.

The precoder enters the surface integral only through the element weights $x_{j(n)}$, which multiply the incident polarization of each path. Writing $\mathbf{e}_n = x_{j(n)}\, \boldsymbol{\psi}_n$ and using $\tilde{\mathbf{G}}(\mathbf{r}) = \tilde{\boldsymbol{\Psi}}(\mathbf{r})\boldsymbol{\Phi}(\mathbf{r})\mathbf{J}$ from Eq. (14), the whole-body absorbed power splits into a phantom part and a scene part,

$$P_{\text{abs}} = \sum_{n,n'} \mathbf{e}_n^H\, \mathbf{T}(\hat{\boldsymbol{k}}_n, \hat{\boldsymbol{k}}_{n'})\, \mathbf{e}_{n'}, \quad (18)$$

where the absorption kernel is the 3×3 matrix

$$\mathbf{T}(\hat{\boldsymbol{k}}, \hat{\boldsymbol{k}}') = \frac{\sigma}{4\sqrt{\alpha(\hat{\boldsymbol{k}})\,\alpha(\hat{\boldsymbol{k}}')}} \times \int_\Sigma \mathbf{F}(\mathbf{r}; \hat{\boldsymbol{k}})^H\, \mathbf{F}(\mathbf{r}; \hat{\boldsymbol{k}}')\, e^{-ik_0(\hat{\boldsymbol{k}}' - \hat{\boldsymbol{k}})\cdot\mathbf{r}}\, \mathrm{d}A. \quad (19)$$

Here $\mathbf{F}(\mathbf{r};\hat{\boldsymbol{k}})$ is the per-path Fresnel operator of Section III for a wave arriving along $\hat{\boldsymbol{k}}$, and $\alpha(\hat{\boldsymbol{k}})$ its depth attenuation. The kernel depends on the phantom alone. It is independent of the path count and the scene, one fixed object per phantom and frequency, evaluated once and reused for every array and every scene. The diagonal $\mathbf{T}(\hat{\boldsymbol{k}},\hat{\boldsymbol{k}})$ is the body's directional absorptivity along $\hat{\boldsymbol{k}}$. The off-diagonal $\mathbf{T}(\hat{\boldsymbol{k}},\hat{\boldsymbol{k}}')$ is the body's spatial Fourier component at the difference wave vector $\mathbf{q}=k_0(\hat{\boldsymbol{k}}'-\hat{\boldsymbol{k}})$, the coherent coupling that a diagonal reduction discards. Per-phantom kernels are pre-computed once per frequency and made available publicly for each phantom.

A scene samples the kernel at its ray directions. Contracting the kernel with the incident polarizations of one path pair gives the path-pair Gram matrix

$$M_{nn'}=\boldsymbol{\psi}_n^H\,\mathbf{T}(\hat{\boldsymbol{k}}_n,\hat{\boldsymbol{k}}_{n'})\,\boldsymbol{\psi}_{n'},\tag{20}$$

an $N\times N$ Hermitian positive-semidefinite matrix whose diagonal is the single-path absorbed power and whose off-diagonal is the coherent coupling. Grouping the paths by the element that radiates each, through the dispatch $\mathbf{J}\in\{0,1\}^{N\times M}$ of Section II, the operator of Eq. (17) factors as

$$\mathbf{Q}=\mathbf{J}^T\,\mathbf{M}\,\mathbf{J}.\tag{21}$$

The factorization separates the array from the body. The dispatch $\mathbf{J}$ is combinatorial, a $\{0,1\}$ assignment of each path to the element that radiates it. The path-pair Gram matrix $\mathbf{M}$ is the continuous part, the scene's sampling of the fixed body through its path geometry, shape, and tissue, and holds none of the array hardware. A change of array updates only $\mathbf{J}$, whereas a change of scene updates only $\mathbf{M}$. The body enters both only through the pre-computed kernel $\mathbf{T}$ of Eq. (19).

For a surface-only refinement, choose a different per-path Fresnel operator $\mathbf{F}(\mathbf{r})$ in Eq. (19). Any bounded linear $\mathbf{F}(\mathbf{r}):\mathbb{C}^3\to\mathbb{C}^3$ from the incident polarization to the transmitted surface field keeps $\mathbf{Q}=\int_\Sigma\tilde{\mathbf{G}}^H\tilde{\mathbf{G}}\,\mathrm{d}A$ an integral of Gram matrices, so the operator stays positive-semidefinite and the beamformer of Section V applies without change. The factorization $\mathbf{Q}=\mathbf{J}^T\mathbf{M}\mathbf{J}$ and its assembly are unchanged.

Surface-only refinements include shadow-edge diffraction through a Fock transition gate on the Fresnel coefficients, cross-body diffraction, surface curvature, broadband operation, and the tissue model of a chosen phantom. In the uplink at mmWave frequencies, the body remains mostly in the radiative far field of the UE, but the spherical wavefronts require the use of spherical Green functions instead of plane waves. A multilayer stack is different: it replaces the homogeneous scalar depth factor with a vector depth Gram. Exact layered TM forward and backward fields have opposite normal components, so a rank-one reduction must be validated for the chosen stack. The supplementary material lists these refinements.

Above 6 GHz, the local APD restriction applies alongside the whole-body restriction, so its evaluation requires the APD averaged over a small patch. Let $\omega_A(\mathbf{r};\mathbf{r}_0)$ be the normalized averaging mask of the $4\,\mathrm{cm}^2$ patch centered at $\mathbf{r}_0$, the spatial-averaging area of the ICNIRP FR2 basic restriction: $\omega_A=1/A$ inside the patch and zero outside, so that $\int_\Sigma\omega_A\,\mathrm{d}A=1$. The local exposure operator is

$$\mathbf{Q}_A(\mathbf{r}_0)=\int_\Sigma\omega_A(\mathbf{r};\mathbf{r}_0)\,\tilde{\mathbf{G}}(\mathbf{r})^H\,\tilde{\mathbf{G}}(\mathbf{r})\,\mathrm{d}A=\mathbf{J}^T\,\mathbf{M}_{A,\mathbf{r}_0}\,\mathbf{J},\tag{22}$$

with $\mathbf{M}_{A,\mathbf{r}_0}$ the masked path-pair Gram matrix, derived in the supplementary material. The form $\mathbf{x}^H\mathbf{Q}_A(\mathbf{r}_0)\,\mathbf{x}$ is the spatially averaged APD over the patch, in $\mathrm{W/m^2}$. With the unnormalized mask $\omega\equiv 1$, the integral returns the whole-body operator $\mathbf{Q}$. The mask multiplies only the surface integrand, so the local operator has the same factorization. Section V constrains this operator with the local basic restriction.

V. Exposure-Constrained Beamforming (ECBF)

The 2020 ICNIRP guidelines [2] set a whole-body restriction and one or more applicable local peak-absorption restrictions. A compliant BS must satisfy all of them on every exposed body at all times. A BS serves K users with precoders $\{\mathbf{x}_k\}_{k=1}^K$. One ray-traced environment gives a single path set shared across all bodies in the scene, so the dispatch $\mathbf{J}$ and the phase $\boldsymbol{\Phi}$ are common, and each body u has its own Fresnel filter through the per-body operators $\mathbf{Q}^{(u)}=\mathbf{J}^T\mathbf{M}^{(u)}\mathbf{J}$ and $\mathbf{Q}_A^{(u)}(\mathbf{r}_0)=\mathbf{J}^T\mathbf{M}_{A,\mathbf{r}_0}^{(u)}\mathbf{J}$ of Section IV. The ECBF sum-rate problem is

$$\begin{aligned}\max_{\{\mathbf{x}_k\}}\;& R(\{\mathbf{x}_k\})\\ \text{s.t.}\;& \textstyle\sum_k\mathbf{x}_k^H\,\mathbf{Q}^{(u)}\,\mathbf{x}_k\le m_u\,\mathrm{SAR}_{\mathrm{wb,lim}}, && \forall u,\\ & \textstyle\max_{\mathbf{r}_0}\sum_k\mathbf{x}_k^H\,\mathbf{Q}_A^{(u)}(\mathbf{r}_0)\,\mathbf{x}_k\le S_{\mathrm{ab,lim}}, && \forall u,\\ & \textstyle\sum_k\|\mathbf{x}_k\|^2\le P,\end{aligned}\tag{23}$$

with m_u the mass of body u, $\mathrm{SAR}_{\mathrm{wb,lim}}$ and $S_{\mathrm{ab,lim}}$ the whole-body and local basic restrictions, and P the transmit-power limit. The first constraint limits the wbSAR of body u, and the second limits its peak spatially averaged APD for the area represented by $\mathbf{Q}_A$. The second row is written for one local averaging area. Where a second area applies, the same row is repeated with its corresponding operator and limit. Every operator is positive-semidefinite, so each exposure constraint is convex in the precoders.

The general-public local APD limit averaged over $4\,\mathrm{cm}^2$ is $20\,\mathrm{W/m^2}$. Above 30 GHz, the $1\,\mathrm{cm}^2$ average has a limit of $40\,\mathrm{W/m^2}$ and is an additional restriction. A compliant precoder enforces every applicable local restriction over the complete body surface. It adds any violated patch constraint until the complete surface satisfies each limit.

When only one local specification applies, one body of mass m and surface area $|\Sigma|$ has one whole-body and one local restriction, which reduce to one ratio. Both scale linearly with the transmit power $\|\mathbf{x}\|^2$ at fixed precoder shape. The whole-body restriction is reached at the transmit power $P_{\mathrm{tx}}^{\mathrm{wb}}=\|\mathbf{x}\|^2\,\mathrm{SAR}_{\mathrm{wb,lim}}\,m/\mathbf{x}^H\mathbf{Q}\mathbf{x}$, and the local restriction at $P_{\mathrm{tx}}^{\mathrm{loc}} = \|\mathbf{x}\|^2 S_{\mathrm{ab,lim}} / \max_{\mathbf{r}_0} \mathbf{x}^H \mathbf{Q}_A(\mathbf{r}_0)\mathbf{x}$. Their ratio is the *binding ratio* $R_{\mathrm{bind}}^{\mathrm{wb>loc}}$,

$$R_{\mathrm{bind}}^{\mathrm{wb>loc}} = \frac{P_{\mathrm{tx}}^{\mathrm{wb}}}{P_{\mathrm{tx}}^{\mathrm{loc}}} = \frac{\eta_A(\mathbf{x})}{\eta_{\mathrm{crit}}}, \quad \eta_A(\mathbf{x}) = |\Sigma| \frac{\max_{\mathbf{r}_0} \mathbf{x}^H \mathbf{Q}_A \mathbf{x}}{\mathbf{x}^H \mathbf{Q} \mathbf{x}}, \tag{24}$$

where the norm factors cancel and the second equality follows by substitution, with $\eta_A(\mathbf{x})$ the precoder's peak-to-mean contrast on the body and $\eta_{\mathrm{crit}} = (S_{\mathrm{ab,lim}}/\mathrm{SAR}_{\mathrm{wb,lim}})\,|\Sigma|/m$ the body threshold. $R_{\mathrm{bind}}^{\mathrm{wb>loc}} > 1$ means the local restriction is reached before the whole-body restriction as transmit power increases. The ratio $S_{\mathrm{ab,lim}}/\mathrm{SAR}_{\mathrm{wb,lim}}$ is $250\,\mathrm{kg/m^2}$ for both the occupational (100/0.4) and the general-public (20/0.08) limits, so $\eta_{\mathrm{crit}} = 250\,|\Sigma|/m$. Above 30 GHz no single binding ratio describes compliance. The first-binding state is instead the smallest permitted transmit-power scale among wbSAR and the two local averages. Since $|\Sigma|/m$ scales as $m^{-1/3}$, a smaller body has a larger threshold. Section VI measures the predicted crossover.

The general problem has no closed form because its constraints must hold for every exposed body u and every applicable patch center $\mathbf{r}_0$. With one body, one active exposure constraint, and one data stream, the design maximizes received signal power subject to that constraint and the power limit,

$$\max_{\mathbf{x}} |\mathbf{h}^T \mathbf{x}|^2 \;\; \text{s.t.} \;\; \mathbf{x}^H \mathbf{Q}\,\mathbf{x} \le m\,\mathrm{SAR}_{\mathrm{wb,lim}}, \;\; \|\mathbf{x}\|^2 \le P, \tag{25}$$

written here for the whole-body restriction. The local case replaces $\mathbf{Q}$ by $\mathbf{Q}_A(\mathbf{r}_0)$ and the limit by $S_{\mathrm{ab,lim}}$. The objective is the squared modulus of a linear functional of $\mathbf{x}$, and both constraints are convex quadratic forms, so this is a non-convex Quadratically Constrained Quadratic Program (QCQP) with two positive-semidefinite constraints. Strong duality holds for this class [7], so the stationarity condition below characterizes the global optimum rather than a local one.

Let $\mathbf{B} = \lambda\mathbf{Q} + \nu\mathbf{I}$. When $\mathbf{B}$ is nonsingular, stationarity gives the maximizer of Eq. (25) in closed form,

$$\mathbf{x}^\star = c^\star \frac{\mathbf{B}^{-1}\mathbf{h}^*}{\|\mathbf{B}^{-1}\mathbf{h}^*\|}, \tag{26}$$

with $0 < c^\star \le \sqrt{P}$ the largest feasible amplitude and Lagrange multipliers $\lambda \ge 0$ and $\nu \ge 0$ fixed by complementary slackness. At least one constraint is active: $c^\star = \sqrt{P}$ when the power constraint binds, while $c^\star < \sqrt{P}$ can occur when the exposure limit binds first. This is the regularizer-shifted matched filter of Ying et al. [7], derived there for the FDTD-computed SAR matrix and now using the analytical operator $\mathbf{Q}$ of Eq. (17) in place of the simulated one. Up to a physically irrelevant common phase, the solution is unique when $\mathbf{B}$ is nonsingular. If $\mathbf{B}$ is singular, its ordinary inverse does not exist. Replacing $\mathbf{B}^{-1}$ by its Moore–Penrose inverse $\mathbf{B}^{+}$ selects the minimum-norm solution; other optima can differ by a null-space component that changes neither the absorbed power nor the served signal. The form needs only $\mathbf{Q} \succeq 0$, which holds for the exact operator and every refinement of it. The only computation is one $M \times M$ element-space solve in Eq. (26). As an array processor, this precoder is a diagonally loaded Minimum-Variance Distortionless Response (MVDR) filter in which the body's absorption covariance replaces the interference covariance.

The two limits of the exposure constraint are known precoders (Fig. 3). A slack constraint ($\lambda \to 0$) gives $\mathbf{x}^\star \propto \mathbf{h}^*$, MRT. A binding constraint ($\lambda \to \infty$ at finite ν) gives $\mathbf{x}^\star \approx \mathbf{Q}^{-1}\mathbf{h}^*$ when $\mathbf{Q}$ is full rank, and the projection of $\mathbf{h}^*$ onto the null space of $\mathbf{Q}$ otherwise, steering the precoder into the absorption null space of the body. Without the power limit, maximizing signal per absorbed power gives the Pareto-optimal precoder $\mathbf{x}^\star \propto \mathbf{Q}^{-1}\mathbf{h}^*$. The constrained design reduces absorption without a proportional loss of signal because the signal and the body combine the same paths differently, one branch per channel in Fig. 3. First, the signal branch under MRT superposes scalar projections $h_j^* h_j \in \mathbb{R}_+$ at one point, whereas the body field superposes vector contributions in $\mathbb{C}^3$, so matched path magnitudes can cancel at $\mathbf{r}$ when their transmitted polarizations differ. Second, the signal channel projects each path through a scalar antenna inner product, whereas the exposure channel applies a 3×3 rank-two Fresnel operator tied to the surface normal. A grazing path still serves the signal but contributes little to absorption, because both Fresnel coefficients vanish at $\mu = 0$. Third, the signal is evaluated at the single UE point, where MRT aligns the path phases. Exposure is evaluated at every $\mathbf{r} \in \Sigma$, where each path has the extra phase $e^{-ik_0 \hat{\boldsymbol{k}}_n \cdot (\mathbf{r} - \mathbf{r}_{\mathrm{UE}})}$ relative to the UE, so alignment at the UE need not hold on the body.

VI. Numerical demonstration

The numerical results use the ray-traced indoor factory shown in Fig. 1. The full Cartesian product spans four phantoms, three URA sizes, LOS and NLOS, 16 scene realizations (each with differently placed concrete scatterers), eight frequencies from 8 to 100 GHz, and nine served-user positions from 6 to 22 m. The first block evaluates the MRT precoder at 21 body-relative device placements and supplies the absorbed-power and binding-restriction results. The second block evaluates six whole-body ECBF limits for the chest focus point at skin contact and supplies the signal–exposure trade-off results. The focus point is the location of the dipole antenna of the UE. Incoherent dosimetry sums powers and therefore cannot determine which basic restriction binds first as transmit power increases, where peak absorption occurs, or how strongly a coherent precoder can concentrate absorption. As previewed in the insets of Fig. 1, the exposure-constrained precoder of Section V then demonstrates its control over coherent concentration.

A. Configuration

Fig. 1 shows the $40 \times 20 \times 5$ m factory scene, which follows the exposure scenarios of Shikhantsov et al. and Wydaeghe et al. [3], [18]. Its Poisson-distributed metallic clutter is inspired by the Indoor Factory with Sparse Clutter and a High Base Station (InF-SH) model of Third Generation

Partnership Project (3GPP) Technical Report 38.901 [19]. A URA of isotropic, vertically polarized elements hangs at one end of the hall at 3 m height, with a 10° mechanical downtilt towards the floor. The exposed body stands in the corridor in front of the array. Sionna RT [12] ray-traces each scene once at the array phase center, using 10^6 ray samples and up to five reflection bounces. The departure direction then gives the analytical phase at every physical element of the square, half-wavelength URA [20]. This construction preserves the intended 10° panel tilt for all three array sizes. Two propagation conditions were ray-traced: LOS and NLOS, with the latter created by a $0.5 \times 2 \times 5\,\mathrm{m}^3$ metallic blocker on the direct path.

Table I lists the phantoms and all the different sweep variables. It gives the average along each axis of the two tabled quantities under MRT: wbSAR in mW/kg and the peak averaged APD $S^{4\,\mathrm{cm}^2}_{\mathrm{ab,max}}$ in $\mathrm{mW/m}^2$, both per watt of conducted power. The calculation uses 2 304 ray-traced scenes and four phantoms, giving 9 216 body–ray combinations. Analytical expansion to 16, 64, and 256 elements gives 27 648 body–scene–array cases. The full Cartesian product spans four phantoms, both propagation conditions, 16 scene realizations, eight frequencies from 8 to 100 GHz, and nine corridor positions from 6 to 22 m of BS-to-UE range. Each configuration uses its own frequency and an array with half-wavelength element spacing at that frequency. Thus the frequency sweep changes the propagation and aperture as well as the tissue response.

The chest, head, and hip anchor triangles lie along the phantom's sagittal plane at heights equal to 0.62, 0.92, and 0.52 of its body height, respectively. At each anchor, the served point moves along the triangle's outward surface normal through device separations of 0, 1, 2, 3, 5, 10, and 20 cm. The MRT block evaluates all 21 placements, giving 580 608 full-body maps. The separate ECBF block evaluates six whole-body absorbed-power limits for the chest focus point at contact, giving 165 888 body-maps. This totals 746 496 body-maps. These placement and precoder sweeps require no additional ray tracing. The exposure operator depends only on the phantom, propagation condition, scene realization, frequency, array, and corridor position. Thus, the device separation and precoder enter only through the user channel $\mathbf{h}$.

The full sweep is computationally feasible because each body–direction response is built once and reused across array sizes, placements, and precoders. An FDTD simulation of this scene requires more than 10^{12} cells for one precoder, one body, and one position. Each new precoder or position requires another simulation. The exposure operator is instead assembled from the ray-traced paths. The complete 746 496-body-map calculation completed in 2.2 hours on one NVIDIA A6000 GPU, or about 10 ms per body-map. Prior full-wave dosimetry of a comparable factory hall covered the head region of one phantom, tens of directions of arrival, and ten environment realizations on a high-performance cluster [18], [20].

TABLE I. Sweep axes and their marginal exposure means under the MRT precoder only. Frequency-by-array entries give wbSAR / $S^{4\,\mathrm{cm}^2}_{\mathrm{ab,max}}$ [$S^{1\,\mathrm{cm}^2}_{\mathrm{ab,max}}$] in mW/kg / $\mathrm{mW/m}^2$ per conducted watt. Bold marks each subtable's maximum. All 580 608 MRT body-maps enter the marginals.

(a) Phantom† [21]		η_{crit}	wbSAR	$S^{4\,\mathrm{cm}^2}_{\mathrm{ab,max}}$
Thelonious	M6	11.3	**0.126**	21.9
Eartha	F8	8.9	0.098	22.0
Ella	F26	6.9	0.073	**23.0**
Duke	M34	6.5	0.069	**23.0**

(b) Freq. × M	16	64	256
8 GHz	0.015 / 4.0	0.053 / 13.0	0.201 / 48.1
10 GHz	0.018 / 5.6	0.050 / 14.5	0.166 / 46.4
12 GHz	0.016 / 4.9	0.051 / 13.6	0.184 / 47.5
15 GHz	0.017 / 5.0	0.051 / 13.8	0.179 / 45.5
20 GHz	0.017 / 4.4	0.053 / 12.8	0.193 / 44.5
28 GHz	0.018 / 4.5	0.055 / 13.3	0.200 / 46.9
60 GHz	0.021 / 5.1 [7.2]	0.064 / 15.2 [20.0]	0.230 / 52.5 [66.7]
100 GHz	0.024 / 5.7 [8.0]	0.071 / 16.2 [21.4]	**0.248** / **56.0** [**71.2**]

(c) Range × condition	LOS	NLOS
6 m	**0.505** / **121.7**	0.049 / 8.6
8 m	0.272 / 68.5	0.060 / 10.8
10 m	0.168 / 45.0	0.049 / 8.9
12 m	0.114 / 29.7	0.038 / 7.2
14 m	0.079 / 21.2	0.035 / 6.3
16 m	0.059 / 16.6	0.026 / 5.3
18 m	0.050 / 15.0	0.029 / 6.5
20 m	0.038 / 12.6	0.019 / 4.6
22 m	0.034 / 10.9	0.022 / 4.8

(d) Separation × anchor	Chest	Head	Hip
0 cm	**15.1**	11.0	11.2
1 cm	14.2	10.4	10.3
2 cm	13.2	9.4	9.7
3 cm	12.5	8.7	9.1
5 cm	11.5	8.3	8.5
10 cm	10.9	7.9	7.8
20 cm	10.8	7.6	7.9

†IT'IS Virtual Population v3.x phantoms, coded by sex and age in years. $\eta_{\mathrm{crit}} = 250|\Sigma|/m$ is the contrast threshold of Eq. (24).

B. *Absorbed power*

The available compliance margin is quantified first. The Equivalent Isotropic Radiated Power (EIRP) is set to 55 dBm, representative of a ceiling-mounted local-area BS. For the 256-element array, 316 W is the equivalent isotropically radiated power. Division by the ideal array gain gives 1.2 W of conducted power. Under MRT, the worst configuration reaches 4.6% of the applicable local restriction and 3.5% of the whole-body restriction. The median configuration reaches 0.3% of its first applicable restriction. Because both quantities scale linearly with transmit power, the EIRP at which each configuration reaches a restriction does not depend on the chosen reference EIRP. Across the sweep, the local restriction is reached at a median 81.1 dBm and a worst-case 68.4 dBm. The child's whole-body average reaches its restriction at a worst-case 69.6 dBm.

Along the corridor, absorbed power depends on the channel geometry. At 6 m, the LOS means are 0.505 mW/kg and 121.7 $\mathrm{mW/m}^2$ per conducted watt, respectively. By 22 m they decrease to 0.034 mW/kg and 10.9 $\mathrm{mW/m}^2$. The NLOS means vary much less over the same range, from 0.049 to 0.022 mW/kg and from 8.6 to 4.8 $\mathrm{mW/m}^2$. Thus, absorbed power depends on both the propagation condition and the BS-to-UE distance.

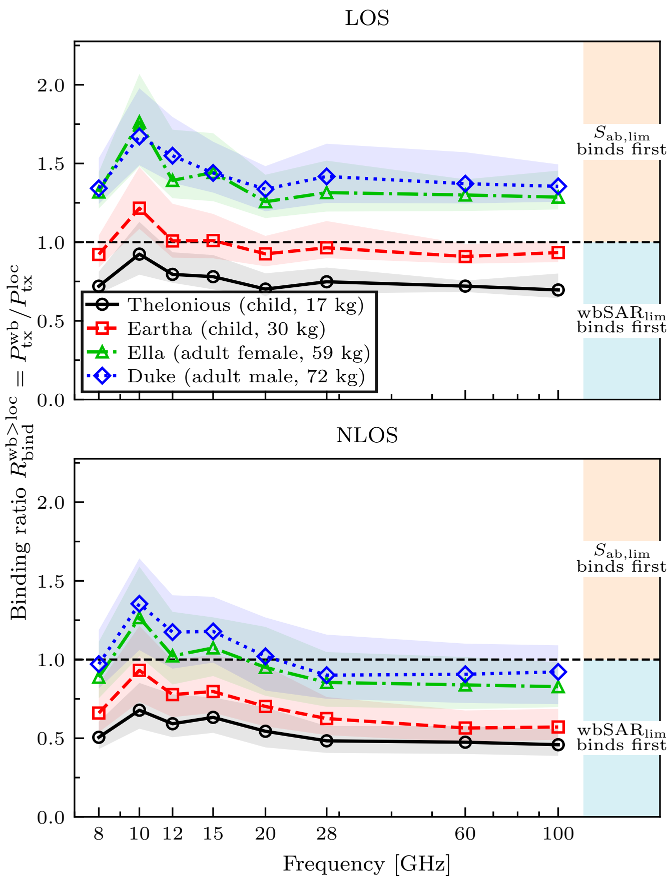

Fig. 4. Binding ratio $R_{\mathrm{bind}}^{\mathrm{wb>loc}}$ against frequency for the four phantoms under LOS (top) and NLOS (bottom). The ratio compares the transmit powers allowed by wbSAR and the applicable local limit. Above the dashed line, a local restriction binds first. Above 30 GHz, the local limit is the lower of the simultaneous $1\,\mathrm{cm}^2$ and $4\,\mathrm{cm}^2$ limits. The lines are the medians over the remaining axes, with interquartile bands.

Moving the focus point away from the surface changes the local value. At the chest, the mean $4\,\mathrm{cm}^2$ value at the served point decreases from $15.1\,\mathrm{mW/m}^2$ at contact to $10.8\,\mathrm{mW/m}^2$ at 20 cm per conducted watt. The corresponding head and hip values decrease from 11.0 and 11.2 to 7.6 and $7.9\,\mathrm{mW/m}^2$. All 21 placements reuse the same exposure operator and differ only through the user channel.

C. Binding restriction

Fig. 4 shows which restriction binds first across frequency and phantom. Below 30 GHz, the binding ratio $R_{\mathrm{bind}}^{\mathrm{wb>loc}} = \eta_A/\eta_{\mathrm{crit}}$ of Eq. (24) is the transmit power at the whole-body limit on wbSAR divided by the power at the local limit on the $4\,\mathrm{cm}^2$ average $S_{\mathrm{ab,max}}^{4\,\mathrm{cm}^2}$. $R_{\mathrm{bind}}^{\mathrm{wb>loc}}$ compares the transmit power allowed by each restriction: above one, the local restriction is reached before the whole-body restriction as power increases. The local restriction binds first in 99.0% of the adult LOS evaluations and 63.4% of the adult NLOS evaluations. For the two children, the corresponding fractions are 38.2% and 13.5%. Their lower mass makes the whole-body average the lower limit in more cases. The threshold $\eta_{\mathrm{crit}} = 250|\Sigma|/m$ ranges from 6.5 for the 72 kg adult to 11.3 for the 17 kg child (Table I).

Above 30 GHz, the $1\,\mathrm{cm}^2$ restriction is additional, so the local limit in Fig. 4 is the lower of the two local limits. Across 145 152 high-band MRT placement evaluations, wbSAR binds first in 62.08%, the $4\,\mathrm{cm}^2$ average in 37.19%, and the $1\,\mathrm{cm}^2$ average in 0.73%. Across the 41 472 high-band ECBF evaluations, the corresponding fractions are 57.86%, 31.78%, and 10.36%. Thus, one binding ratio is not a complete high-band compliance description. A compliant optimizer retains all three constraints.

The averaging-area sweep shows how spatial averaging changes the peak for a fixed precoder. At 28 GHz, the peak over $1\,\mathrm{cm}^2$ is 1.44 to 1.96× its $4\,\mathrm{cm}^2$ value across the four phantoms, while the $20\,\mathrm{cm}^2$ peak is 0.45 to 0.53× that value (Fig. 5a). Thus, the hypothetical binding ratio increases as the averaging area decreases (Fig. 5b). At 1 and $4\,\mathrm{cm}^2$, all four phantoms are above $R = 1$. At $20\,\mathrm{cm}^2$, none are above it. The $4\,\mathrm{cm}^2$ values range from 1.14 to 2.19. This sweep holds the local limit fixed to isolate spatial averaging. In the regulatory comparison, the $1\,\mathrm{cm}^2$ limit is twice the $4\,\mathrm{cm}^2$ limit.

The ratio of the single-triangle peak to the $4\,\mathrm{cm}^2$ average also depends on frequency (Fig. 5c). It is 1.19 to 1.29 at 8 GHz, 2.49 to 3.06 at 28 GHz, and reaches 3.75 to 4.80 at 60 GHz before decreasing at 100 GHz. Below 30 GHz only the $4\,\mathrm{cm}^2$ restriction applies. Above 30 GHz, complete compliance evaluates the $1\,\mathrm{cm}^2$ and $4\,\mathrm{cm}^2$ body-maps separately with their respective limits. The binding restriction therefore depends on the propagation condition and averaging area.

D. Coherent concentration

Fig. 6 shows the peak-to-mean contrast for four precoders, the change in on-body peak with device separation, and the corresponding absorption-mode loading. The concentration is measured by the peak-to-mean contrast

$$\eta(\mathbf{x}) = |\Sigma| \max_{\mathbf{r}\in\Sigma} S_{\mathrm{ab}}(\mathbf{r}) \Big/ \int_{\Sigma} S_{\mathrm{ab}}\, \mathrm{d}A, \tag{27}$$

the local peak per watt absorbed. The absorbed power scales both the numerator and denominator. Thus η is independent of array gain and transmit power. The metric measures how unevenly a given amount of absorbed power is distributed over the body.

Table II lists how four considered precoders affect the peak-to-mean contrast η. The incoherent-multipath precoder preserves the array pattern but sums the arriving paths in power, giving a contrast of 3.30 under LOS and 2.02 under NLOS. The corresponding MRT values are 7.17 and 6.75. Skin-aimed ECBF, defined as ECBF focused at a point on the skin under the strictest absorbed-power limit, reaches 35.56 and 24.82. These values are medians over four Duke scene realizations at 28 GHz and 256 elements. The worst case over all precoders is larger. Appendix A shows that it depends on the retained rank.

To separate phase coherence from array gain, the same illumination is replayed with the phase relation between

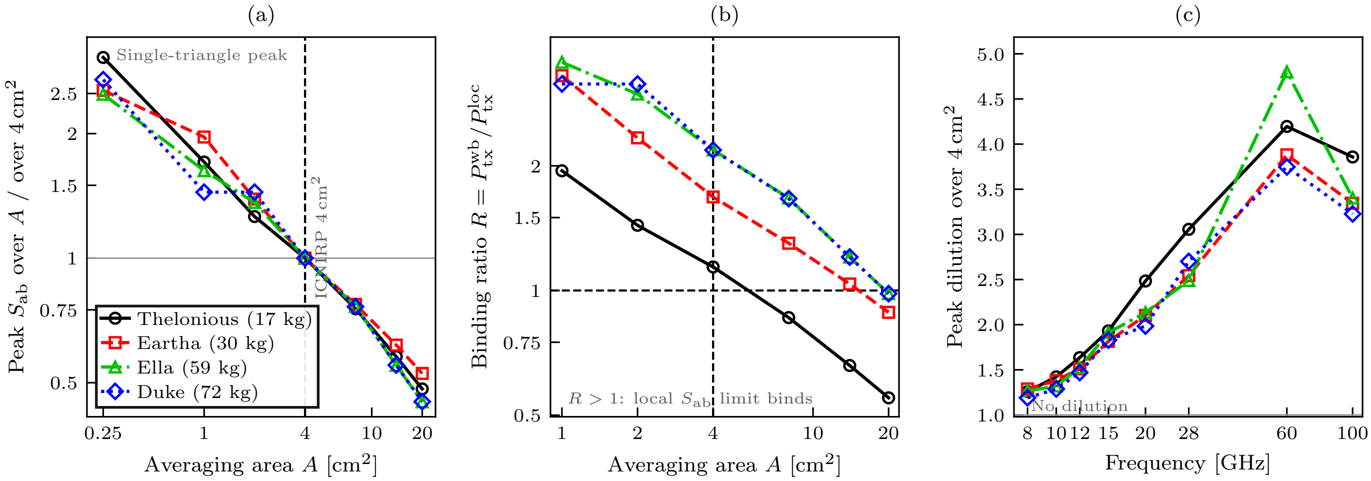

Fig. 5. Averaging area changes the binding restriction for a fixed precoder. (a) Peak APD, normalized by its $4\,\text{cm}^2$ value at 28 GHz. (b) Hypothetical binding ratio R against averaging area when the local limit is held fixed. (c) Ratio of the single-triangle peak to the $4\,\text{cm}^2$ average under LOS, extended through 60 and 100 GHz. Above 30 GHz, complete regulatory compliance instead enforces the distinct $1\,\text{cm}^2$ and $4\,\text{cm}^2$ limits simultaneously.

TABLE II. The impact of precoder and device separation on the median peak-to-mean contrast η of (27) across four Duke scene realizations at 28 GHz and 256 elements

Precoder	Focused at	LOS	NLOS	Separation-dependent
Incoherent multipath	None	3.30	2.02	No
Scrambled arrivals	None	6.60	6.46	No
MRT	UE	7.17	6.75	No
Skin-aimed ECBF	Skin point	35.56	24.82	Yes
Worst case	None	See Appendix A		No

arrivals randomized. The replay produces a speckle-like field. The peak-to-mean contrast is 6.60 under LOS and 6.46 under NLOS. On the same channels, MRT reaches 7.17 and 6.75. Thus, phase alignment at the user and the hot spot on the body are unrelated. A hot spot requires a precoder focused on the skin and sufficient aperture. Skin-aimed ECBF exceeds the scrambled result by 5.39× under LOS and 3.84× under NLOS.

Moving the served device away from the skin changes the two main precoders differently (Fig. 6b). The MRT hot spot is generally not at the device. At contact, the median ratio between the value at the served point and the body maximum is 0.89 under LOS and 0.40 under NLOS. At 20 cm, the ratios are 0.51 and 0.26. The served-device location thus does not determine the location of peak absorption.

Skin-aimed ECBF loses its concentration with separation. At contact, almost all of its peak is at the focus point. The focus-point-to-body-peak ratio falls to 0.39 under LOS and 0.23 under NLOS at 2 cm, and reaches 0.10 and 0.06 at 20 cm. Fig. 6b shows this decrease. Thus, this precoder concentrates absorption only when focused close to the skin. The MRT precoder does not have this concentration.

Strong concentration requires low total absorption, so the precoder weights directions that the body absorbs weakly. The MRT precoder places 99.6% of its absorbed power in the eight strongest absorption modes. Skin-aimed ECBF places 35.1% there and the other 64.9% in the weakly absorbed tail (Fig. 6c). This mode loading allows the skin-aimed precoder to produce a high local peak while keeping the total absorbed power low. Equation (27) divides a local peak by that total, and the weakly absorbed tail keeps the denominator small.

No link precoder reaches the worst case over all precoders. The worst-case contrast increases with the number of directions the scene provides (Appendix A). Thus, a worst-case value is meaningful only together with its retained rank k.

E. Signal and exposure trade-off

Fig. 7 shows the Pareto front between received signal power and whole-body absorbed power. Varying the absorbed-power limit moves along this front from MRT through ECBF to the exposure-optimal Generalized Eigenvalue (GEV) precoder. At half the MRT signal power, the median absorbed power decreases by a factor of 5.42 under LOS and 5.18 under NLOS. Over the limit sweep, the absorbed power never exceeds the absorbed power under MRT. The plotted configuration reaches 6.27× at this point, compared with the 5.42× LOS sweep median. The GEV optimum is many orders of magnitude below the MRT value. The optimum reduces absorption to the numerical floor at the cost of eliminating the link entirely.

A scalar-limit design cannot follow this front because the ray-traced operator is not diagonal. The pointwise reduction of Section IV permits only a per-element rescaling of the matched filter. Across the four-decade limit sweep, this diagonal-only model remains at the MRT corner. Its absorbed power changes by less than one part in 10^{12} in the plotted replay. Thus, the off-diagonal coupling that the diagonal-only model discards causes the Pareto gain.

The Pareto trade-off also appears spatially as a large change in level with little change in shape. Fig. 8 shows

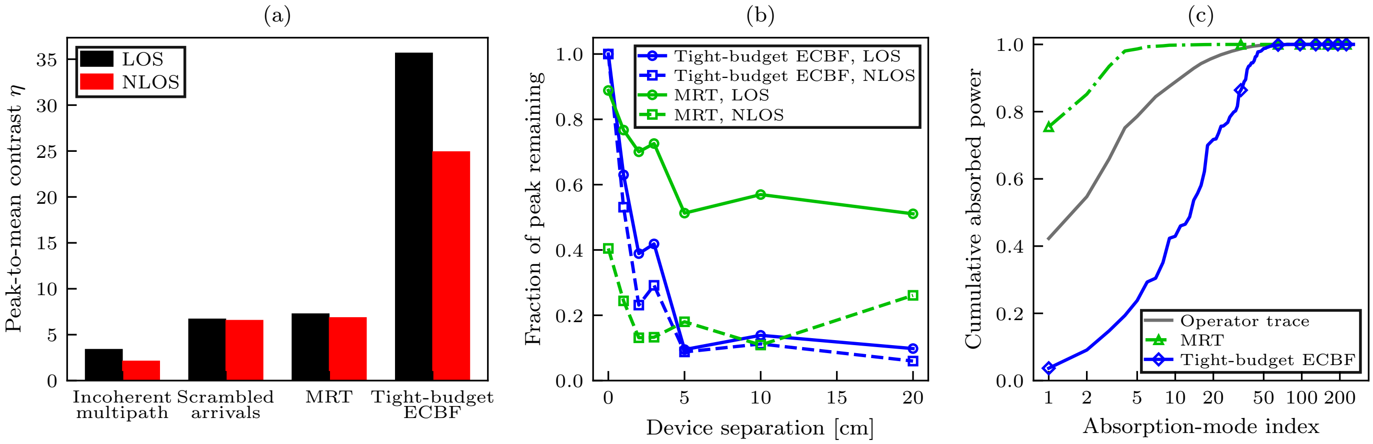

Fig. 6. Peak-to-mean contrast η, on-body peak against device separation, and absorption-mode loading for the four precoders. (a) Median contrast under LOS and NLOS. (b) Median focus-point-to-body-peak ratio for MRT and skin-aimed ECBF as the served point moves away from the skin. Panels (a) and (b) use four Duke scene realizations at 28 GHz and 256 elements. (c) Cumulative absorbed power by absorption-mode index for one of those LOS scenes.

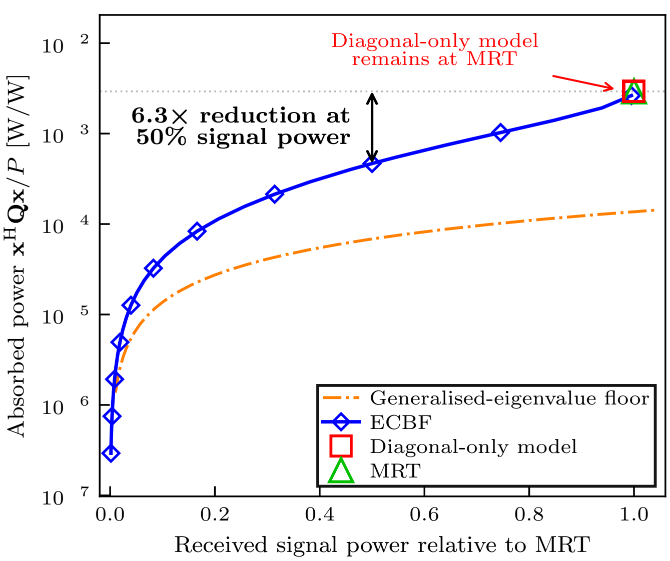

Fig. 7. The signal–exposure Pareto front for one adult configuration. The ECBF design lowers absorbed power by 6.27× at half the MRT signal power in this configuration. The sweep medians are 5.42× under LOS and 5.18× under NLOS. The front approaches the generalized-eigenvalue floor as the signal power decreases, whereas the diagonal-only model remains at the MRT corner.

the absorbed power on an adult body as the absorbed-power limit decreases. The device is at the chest. The MRT precoder produces interference fringes across the whole body. Decreasing the limit to 1% lowers the single-triangle peak by about 78×. The lower limit reduces the hot footprint, defined as the surface above half of each panel's peak, from 849 to 127 cm^2. The peak-to-mean contrast changes from 8.1 to 10.4. Thus, the design lowers the absorbed-power level without concentrating the remaining deposition. The worst-case precoder has a 1.06 cm^2 footprint at a contrast of 65.7 and has no link signal.

The exposure reduction along the Pareto front invites two possibilities. At a fixed compliance limit, a link can use the freed absorption margin for more received signal power. At a fixed received signal power, the body absorbs 5.29× less at the population median. This reduction is an ALARA control for the high-EIRP occupational settings of European Union directive 2013/35/EU [22]. A larger array gives stronger control. The median suppression at half the MRT signal power increases from 4.25× at $M = 16$ to 5.35× at $M = 64$ and 6.45× at $M = 256$. At fixed receiver noise and interference, this half-signal point corresponds to a 3.01 dB SNR reduction. Its throughput and bit-error-rate consequences depend on the selected modulation, coding, and link adaptation.

The whole-body constraint does not cover the local restrictions. Above 30 GHz, a compliant design uses the all-patch formulation of Section V to enforce the whole-body, 4 cm^2, and 1 cm^2 constraints simultaneously. In the campaign, the 1 cm^2 restriction binds first in 10.36% of the high-band ECBF evaluations even though it does so in only 0.73% of the corresponding MRT evaluations.

VII. Conclusion

This work derived a coherent downlink exposure operator and beamformer in closed form from the same ray-traced paths that determine the communication channel. The body reduces to a reusable absorption kernel, and the whole-body absorbed power follows in closed form from the path-pair Gram matrix. An indoor-factory study in FR2 and FR3 tested the construction at a scale beyond full-wave dosimetry, over four phantoms, three square array sizes, eight frequencies from 8 to 100 GHz, two propagation conditions, 16 scene realizations, and nine served-user positions. The MRT and ECBF studies contain 746 496 full-body maps. At half the MRT signal power, the ECBF design suppresses the median whole-body absorbed power by a factor of 5.29, with an interquartile range of 4.43–6.49.

The practical consequence is that exposure-aware beamforming becomes an inner-loop step with no full-wave simulation. Because the operator is built from the ray

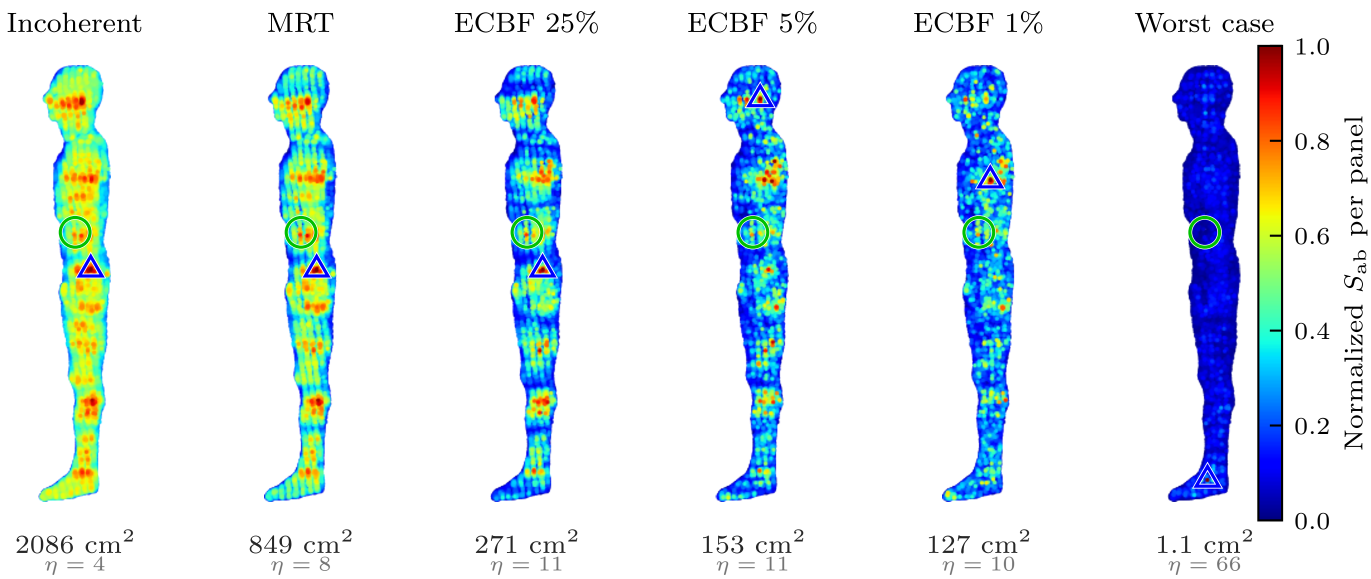

Fig. 8. Normalized APD on an adult body under LOS at 28 GHz, with the device at the chest. The panels compare incoherent multipath, MRT, ECBF limits, and the worst-case precoder. Each panel is normalized by its own peak. The numbers under each panel give the hot-footprint area (the surface above half that panel's peak) and the peak-to-mean contrast η. A circle marks the served point and a triangle marks the peak.

tracer that already estimates the channel, all applicable ICNIRP basic restrictions are evaluated from the same paths. The precoder is re-solved per user as a person moves across the floor. Above 30 GHz the binding restriction is not fixed: wbSAR, the 4 cm^2 average, and the 1 cm^2 average each bind first in part of the campaign. The exposure problem must retain all three. This is a distinction the coherent operator captures and incoherent dosimetry, which superposes powers, cannot. Thus, the method fits well in 5G and 6G compliance under the ALARA principle, and its incoherent limit is the surface dosimetry law of Wydaeghe et al. [11].

Future work includes two directions. The construction is differentiable in every input, so gradients of the regulated quantities propagate to antenna positions, orientations, and beam codebooks, which opens exposure-aware design beyond the precoder. The single-user beamformer will also be extended to multi-body ECBF across an occupied factory floor, with one per-body operator on the shared paths.

Appendix A
Worst case over precoders

Section VI-D uses the adversarial maximum over precoders, derived here as a function of retained rank. It lies at the opposite end of the GEV pencil from the exposure-optimal precoder in Section V.

Define the single-point operator as $\mathbf{P}(\mathbf{r}) = \tilde{\mathbf{G}}(\mathbf{r})^H \tilde{\mathbf{G}}(\mathbf{r})$. Then $S_{\mathrm{ab}}(\mathbf{r}) = \mathbf{x}^H \mathbf{P}(\mathbf{r}) \mathbf{x}$ and $\mathbf{Q} = \int_\Sigma \mathbf{P}\, \mathrm{d}A$. Let $\mathbf{Q}_k$ retain the k largest eigenvalues of $\mathbf{Q}$, and let $\mathbf{Q}_k^+$ denote its pseudoinverse. The focusing density $c(\mathbf{r}) = \mathrm{tr}[\mathbf{Q}_k^+ \mathbf{P}(\mathbf{r})]$ integrates to the retained rank:

$$\int_\Sigma c(\mathbf{r})\, \mathrm{d}A = \mathrm{tr}\big[\mathbf{Q}_k^+ \mathbf{Q}\big] = k. \tag{28}$$

The equality follows because $\mathbf{Q}_k^+$ annihilates the discarded subspace. The total focusing density is exactly k, regardless of its distribution over the surface. Thus, a precoder that concentrates c at one point reduces it elsewhere.

Restrict the precoder to the k retained directions. Whitening by $\mathbf{Q}_k$ reduces Eq. (27) to a Rayleigh quotient. Maximising over the precoder and surface gives

$$\eta_{\max}(k) = |\Sigma| \max_{\mathbf{r}\in\Sigma} \lambda_{\max}\Big[\mathbf{Q}_k^{-1/2} \mathbf{P}(\mathbf{r})\, \mathbf{Q}_k^{-1/2}\Big]. \tag{29}$$

The whitened operator has trace $c(\mathbf{r})$ and rank at most three, with one dimension per field component. Thus, its largest eigenvalue lies between $c(\mathbf{r})/3$ and $c(\mathbf{r})$. Equation (28) gives c a mean of $k/|\Sigma|$:

$$\tfrac{1}{3}|\Sigma| \max_{\mathbf{r}} c \;\le\; \eta_{\max}(k) \;\le\; |\Sigma| \max_{\mathbf{r}} c, \qquad \eta_{\max}(k) \ge \tfrac{k}{3}. \tag{30}$$

The worst case grows at least linearly with the number of retained directions. Averaging over a local-restriction area cannot increase this maximum because each averaged operator is a convex combination of pointwise operators.

The retained rank affects both the reported value and its physical interpretation. It counts the directions that the scene provides. Thus, the adversarial limit is incomplete without k and depends on the hall. For the same body, a scene with more directions has a larger focusing density and a higher limit.

References

[1] T. S. Rappaport *et al.*, "Millimeter wave mobile communications for 5G cellular: It will work!" *IEEE Access*, vol. 1, pp. 335–349, 2013.

[2] ICNIRP, "Guidelines for limiting exposure to electromagnetic fields (100 kHz to 300 GHz)," *Health Phys.*, vol. 118, no. 5, pp. 483–524, 2020.

[3] R. Wydaeghe, S. Shikhantsov, E. Tanghe, G. Vermeeren, L. Martens, P. Demeester, and W. Joseph, “Realistic human exposure at 3.5 and 28 GHz for distributed and collocated MaMIMO in indoor environments using hybrid ray-tracing and FDTD,” *IEEE Access*, vol. 10, pp. 130 996–131 004, 2022, doi: 10.1109/ACCESS.2022.3227107.

[4] R. Wydaeghe, S. Shikhantsov, G. Vermeeren, L. Martens, E. Tanghe, and W. Joseph, “Hybrid ray-tracing-QuaDRiGa/FDTD method for realistic 28 GHz exposure with 6G CF-MaMIMO in 3D outdoor environments,” *npj Wireless Technol.*, vol. 2, no. 1, Art. no. 13, 2026.

[5] B. M. Hochwald and D. J. Love, “Minimizing exposure to electromagnetic radiation in portable devices,” in *Proc. Inf. Theory Appl. Workshop (ITA)*, San Diego, CA, 2012, doi: 10.1109/ITA.2012.6181792.

[6] B. M. Hochwald, D. J. Love, S. Yan, P. Fay, and J.-M. Jin, “Incorporating specific absorption rate constraints into wireless signal design,” *IEEE Commun. Mag.*, vol. 52, no. 9, pp. 126–133, 2014.

[7] D. Ying, D. J. Love, and B. M. Hochwald, “Closed-loop precoding and capacity analysis for multiple-antenna wireless systems with user radiation exposure constraints,” *IEEE Trans. Wireless Commun.*, vol. 14, no. 10, pp. 5859–5870, 2015.

[8] A. Ebadi-Shahrivar, P. Fay, D. J. Love, and B. M. Hochwald, “Determining electromagnetic exposure compliance of multi-antenna devices in linear time,” *IEEE Trans. Antennas Propag.*, vol. 67, no. 12, pp. 7585–7596, 2019.

[9] D. Ying, D. J. Love, and B. M. Hochwald, “Sum-rate analysis for multi-user MIMO systems with user exposure constraints,” *IEEE Trans. Wireless Commun.*, vol. 16, no. 11, pp. 7376–7388, 2017.

[10] M. R. Castellanos, Y. Liu, D. J. Love, B. Peleato, J.-M. Jin, and B. M. Hochwald, “Signal-level models of pointwise electromagnetic exposure for millimeter wave communication,” *IEEE Trans. Antennas Propag.*, vol. 68, no. 5, pp. 3963–3977, 2020.

[11] R. Wydaeghe, L. Martens, G. Vermeeren, E. Tanghe, and W. Joseph, “Closed-form absorbed-power dosimetry from 1 to 100 GHz,” *IEEE Trans. Antennas Propag.*, submitted for publication.

[12] J. Hoydis *et al.*, “Sionna RT: Differentiable ray tracing for radio propagation modeling,” arXiv:2303.11103, 2023.

[13] S. Zhukov, A. Iones, and G. Kronin, “An ambient light illumination model,” in *Rendering Techniques ’98 (Proc. Eurographics Workshop on Rendering)*, G. Drettakis and N. Max, Eds. Springer, 1998, pp. 45–56.

[14] H. Landis, “Production-ready global illumination,” *ACM SIGGRAPH 2002 Course Notes*, vol. 16, 2002.

[15] T. Akenine-Möller, E. Haines, and N. Hoffman, *Real-Time Rendering*, 4th ed. A K Peters/CRC Press, 2018.

[16] S. Kodera, K. Taguchi, Y. Diao, T. Kashiwa, and A. Hirata, “Computation of whole-body average SAR in realistic human models from 1 to 100 GHz,” *IEEE Trans. Microw. Theory Tech.*, vol. 72, no. 1, pp. 91–100, 2024.

[17] S. Boyd and L. Vandenberghe, *Convex Optimization.* Cambridge, U.K.: Cambridge Univ. Press, 2004.

[18] S. Shikhantsov, A. Thielens, G. Vermeeren, E. Tanghe, P. Demeester, L. Martens, G. Torfs, and W. Joseph, “Hybrid ray-tracing/FDTD method for human exposure evaluation of a massive MIMO technology in an industrial indoor environment,” *IEEE Access*, vol. 7, pp. 21 020–21 031, 2019, doi: 10.1109/ACCESS.2019.2897921.

[19] 3GPP, “Study on channel model for frequencies from 0.5 to 100 GHz,” 3rd Generation Partnership Project, Tech. Rep. TR 38.901, 2020.

[20] S. Shikhantsov *et al.*, “Massive MIMO propagation modeling with user-induced coupling effects using ray-tracing and FDTD,” *IEEE J. Sel. Areas Commun.*, vol. 38, no. 9, pp. 1955–1963, 2020.

[21] M.-C. Gosselin *et al.*, “Development of a new generation of high-resolution anatomical models for medical device evaluation: the Virtual Population 3.0,” *Phys. Med. Biol.*, vol. 59, no. 18, pp. 5287–5303, 2014.

[22] “Directive 2013/35/EU of the European Parliament and of the Council of 26 June 2013 on the minimum health and safety requirements regarding the exposure of workers to the risks arising from physical agents (electromagnetic fields),” *Off. J. Eur. Union*, L 179, pp. 1–21, Jun. 2013.