RCIP-BASED REDUCED-BASIS MODELING FOR FREQUENCY SWEEPS IN NEARLY TOUCHING PLASMONIC DISKS*

TIANHAO HU† AND ZECHENG GAN†

Abstract. We develop a reduced-basis method for repeated frequency-domain solution of a two-dimensional quasistatic transmission problem for nearly touching plasmonic disks. The high-fidelity discretization combines a Nyström boundary integral method with recursively compressed inverse preconditioning (RCIP), and the reduced state is the RCIP-transformed density on the coarse grid. For a fixed geometry, the parameter-dependent part of the compressed inverse is confined to local gap blocks. Introducing explicit restriction and extension operators for these blocks yields an offline–online decomposition of the reduced operator and of a discrete residual indicator. The resulting online algebra depends on the reduced dimension and the local RCIP block size rather than on the global coarse-grid dimension, although the local RCIP recursion remains parameter dependent. We also formulate target-dependent layer-potential evaluation and a multi-layer local refinement for targets close to a gap. The reported experiments show rapid decay of snapshot singular values and small discrepancies between the reduced and RCIP solutions for fixed two- and three-disk geometries over the sampled frequency band. These computations support the proposed localized reduction strategy; they do not constitute a certified error bound or a geometry-parametric surrogate.

Key words. reduced-basis method recursively compressed inverse preconditioning boundary integral equation plasmonic resonance many-query frequency sweep

MSC codes. 65R20, 65N38, 78M15

1. Introduction.

Closely spaced metallic particles can support strongly localized quasistatic fields in their intervening gaps [18, 20, 22]. These fields are relevant to sensing and surface-enhanced spectroscopy [6, 17], but they are also difficult to resolve numerically because the boundary density develops increasingly fine spatial scales as the gap closes. The present work considers a two-dimensional local-dielectric model, which represents cross-sections of infinitely long cylinders in the quasistatic limit. At subnanometer gaps, nonlocal response and tunneling can be important [7, 24]; the smallest gaps used below should therefore be read as numerical stress tests of the classical model rather than as quantitatively predictive nanostructures.

Boundary integral equations reduce the transmission problem to the particle boundaries, but close interactions still require specialized discretization. Recursively compressed inverse preconditioning (RCIP) combines local refinement with compression onto a coarse Nyström grid [9]. The rapidly varying physical density is recovered locally, whereas the transformed density is represented on the coarse grid. This separation is attractive for repeated solves because the costly gap resolution is confined to small blocks. It does not, by itself, imply that the frequency-parametric solution manifold has a prescribed Kolmogorov width; compressibility must be assessed from the computed snapshots.

Reduced-basis methods replace repeated high-dimensional solves by projection onto a space generated from selected high-fidelity solutions [12, 19, 21]. Reduced models for boundary integral and electromagnetic scattering formulations have been studied in, for example, [8, 13, 16]. The main difficulty here is that the RCIP compressed inverse depends nonaffinely on the material contrast through a local recursion. Applying a projection without exposing this dependence would retain high-dimensional work at each frequency.

*Submitted to the editors August 10, 2026.

†Hong Kong University of Science and Technology (Guangzhou), Guangzhou, Guangdong, China (thu176@connect.hkust-gz.edu.cn, zechenggan@hkust-gz.edu.cn).

This paper makes the following contributions for a fixed geometry and a fixed set of output targets:

- We use the RCIP-transformed coarse-grid density as the reduced state and distinguish snapshot-SVD diagnostics from residual-greedy basis construction.
- We express the parameter-dependent part of the compressed inverse through explicit local restriction and extension matrices. This gives dimensionally consistent reduced-operator assembly involving only precomputed global factors and small local RCIP blocks.
- We derive a complex-valued discrete residual indicator using Hermitian products. The indicator can be evaluated from a precomputed Gram factor and used on a finite training set, but it is not a certified error bound without a stability lower bound.
- We combine the reduced state with target-dependent product integration and a wider local reconstruction, termed multi-layer RCIP (mRCIP), for targets close to a gap.

The numerical section reports reduced-to-RCIP discrepancies for separately trained two- and three-disk configurations. Geometry and target location are not treated as online parameters, and no continuous-model error estimate is claimed.

Section 2 states the quasistatic boundary integral model and its RCIP discretization. Section 3 presents the reduced-basis formulation, local assembly, and residual indicator. Target evaluation and mRCIP are described in sections 4 and 5; numerical results and limitations follow in sections 6 and 7.

2. Problem Formulation and the RCIP Method.

2.1. Quasistatic model for two disks. Consider a two-dimensional system describing the plasmonic interaction between nearly-touching metallic nanoparticles, governed by the following boundary value problem:

$$
\begin{aligned}
&\Delta\Phi(\mathbf{x}) = 0, && \mathbf{x} \in \mathbb{R}^2 \backslash \Gamma \\
&\Phi\left(\mathbf{x}^-\right) = \Phi\left(\mathbf{x}^+\right), && \mathbf{x} \in \Gamma \\
&\epsilon_p \frac{\partial \Phi\left(\mathbf{x}^-\right)}{\partial \mathbf{n}} = \epsilon_s \frac{\partial \Phi\left(\mathbf{x}^+\right)}{\partial \mathbf{n}}, && \mathbf{x} \in \Gamma \\
&\Phi(\mathbf{x}) = -\mathbf{E}_0 \cdot \mathbf{x} + \mathcal{O}\left(\tfrac{1}{|\mathbf{x}|}\right), \quad \text{as} \quad |\mathbf{x}| \to \infty
\end{aligned}
\tag{2.1}
$$

Here $\mathbf{x} = (x, y) \in \mathbb{R}^2$, and Γ is the union of the boundaries of two disjoint disks with radius $r = 30$ nm and separation $d \ll r$. The model is the two-dimensional cross-section of infinitely long cylinders and neglects retardation. The surrounding permittivity is $\epsilon_{\mathrm{s}} = 1$, and the complex particle permittivity is ϵ_{p}. We use the time convention $\mathbf{E}(\mathbf{x}, t) = \Re\{-\nabla\Phi(\mathbf{x}) \exp(-i\omega t)\}$, for which a passive material has $\operatorname{Im} \epsilon_{\mathrm{p}}(\omega) > 0$. A uniform incident field $\mathbf{E}_0 = (E_0, 0)$ is applied along the x-axis. The origin is the midpoint of the gap, and γ_j $(j = 1, 2)$ denotes the closest point on disk j; see Figure 1.

The metallic permittivity ϵ_{p} is parameterized by the frequency ω based on the Drude-Lorentz model:

$$
\epsilon_{\mathrm{p}}(\omega) = 1 - \frac{\omega_0^2}{\omega(\omega + i\gamma_0\omega_0)} - \sum_{j=1}^{N} \frac{f_j\omega_j^2 + i\gamma_j'\omega_j\omega}{\omega^2 - \omega_j^2 + i\gamma_j\omega_j\omega}
\tag{2.2}
$$

The disks are taken to be silver, with empirical fitting parameters from [23]. The sweep is reported in terms of photon energy, $\hbar\omega \in [0.1, 4.5]$ eV. Because the governing

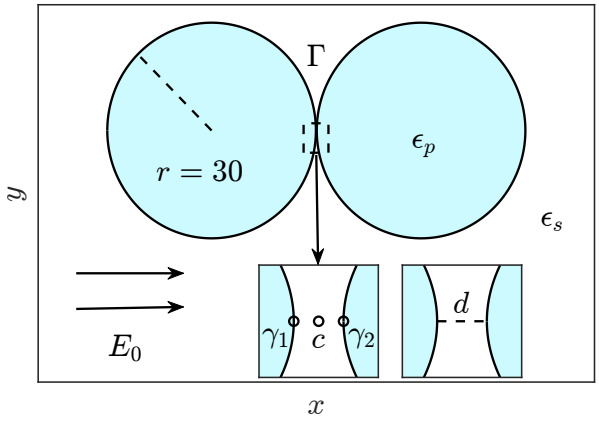

FIG. 1. *Two close disks with radius r and distance d. The permittivity of the disks is* $\epsilon_{\rm p}$ *and the surrounding medium is vacuum* ($\epsilon_{\rm s} = 1$)*. An incident field with intensity* E_0 *is applied from the x-direction.*

equation is quasistatic and two-dimensional, the upper part of this interval should be interpreted within that modeling approximation rather than as a full-wave prediction for finite three-dimensional particles.

(2.1) can be recast into a boundary integral equation (BIE) formulated in terms of an unknown single-layer density $\rho(\mathbf{x})$:

$$\rho(\mathbf{x}) + 2\lambda \int_\Gamma \frac{\partial G(\mathbf{x},\mathbf{y})}{\partial \mathbf{n}_\mathbf{x}} \rho(\mathbf{y}) ds_\mathbf{y} = -2\lambda \frac{\partial \Phi_{\rm in}(\mathbf{x})}{\partial \mathbf{n}}, \quad \mathbf{x} \in \Gamma \tag{2.3}$$

where the material contrast parameter is defined as $\lambda = \frac{\epsilon_p - \epsilon_s}{\epsilon_p + \epsilon_s}$. The 2D free-space Green's function for the Laplace equation is $G(\mathbf{x},\mathbf{y}) = -\frac{1}{2\pi}\log|\mathbf{x}-\mathbf{y}|$, and $\Phi_{\rm in}(\mathbf{x}) = -\mathbf{E}_0 \cdot \mathbf{x}$ represents the incident potential. We can further express the total potential distribution $\Phi(\mathbf{x})$ explicitly in terms of this single-layer potential as:

$$\Phi(\mathbf{x}) = \Phi_{\rm in}(\mathbf{x}) + \int_\Gamma G(\mathbf{x},\mathbf{y})\rho(\mathbf{y}) ds_\mathbf{y} = -\mathbf{E}_0 \cdot \mathbf{x} + \int_\Gamma G(\mathbf{x},\mathbf{y})\rho(\mathbf{y}) ds_\mathbf{y}, \quad \mathbf{x} \notin \Gamma \tag{2.4}$$

The physical quantities of interest include the potential field Φ and the electric field $\mathbf{E}$, which is formulated as:

$$\mathbf{E}(\mathbf{x}) = -\nabla\Phi(\mathbf{x}) = \mathbf{E}_0 - \nabla \int_\Gamma G(\mathbf{x},\mathbf{y})\rho(\mathbf{y}) ds_\mathbf{y}, \quad \mathbf{x} \notin \Gamma \tag{2.5}$$

Since the target evaluation point $\mathbf{x}$ is located away from the boundary Γ, the Green's function $G(\mathbf{x},\mathbf{y})$ exhibits no singularity. Consequently, the integral and the gradient operator can be safely interchanged:

$$\mathbf{E}(\mathbf{x}) = \mathbf{E}_0 - \int_\Gamma \nabla_\mathbf{x} G(\mathbf{x},\mathbf{y})\rho(\mathbf{y}) ds_\mathbf{y}, \quad \mathbf{x} \notin \Gamma \tag{2.6}$$

2.2. Well-posedness and geometric gap singularity. The boundary integral equation (2.3) can be abstracted as the following operator equation:

$$(I + \lambda K)\rho(\mathbf{x}) = \lambda g(\mathbf{x}), \quad \mathbf{x} \in \Gamma, \tag{2.7}$$

where I is the identity operator and the integral operator K acts on the unknown layer density ρ as:

$$(K\rho)(\mathbf{x}) = 2\int_\Gamma \frac{\partial G(\mathbf{x},\mathbf{y})}{\partial \mathbf{n}_\mathbf{x}}\rho(\mathbf{y})ds_\mathbf{y}, \quad \mathbf{x}\in\Gamma. \tag{2.8}$$

With the factor of two in this definition, K is the adjoint Neumann–Poincaré operator in the normalization used here. Its natural setting is a subspace of $H^{-1/2}(\Gamma)$ with the componentwise charge modes removed,

$$\mathcal{H}_0^*(\Gamma) = \left\{ \rho \in H^{-1/2}(\Gamma) : \int_{\Gamma_j} \rho\, ds = 0 \text{ for every boundary component } \Gamma_j \right\}. \tag{2.9}$$

For a fixed positive separation between smooth components, the symmetrized operator is compact and self-adjoint in its energy pairing. After removal of the charge modes, its relevant spectrum is real and contained in $(-1,1)$ under the present normalization [1, 3,4]. This statement depends on the function space and on the sign and normalization conventions for the Green's function.

Invertibility of $I+\lambda K$ requires $-1/\lambda$ to lie outside this spectrum. Recalling

$$\lambda(\omega) = \frac{\epsilon_\mathrm{p}(\omega)-\epsilon_\mathrm{s}}{\epsilon_\mathrm{p}(\omega)+\epsilon_\mathrm{s}}, \tag{2.10}$$

passivity gives

$$\operatorname{Im}\left(-\frac{1}{\lambda(\omega)}\right) = \frac{2\epsilon_\mathrm{s}\operatorname{Im}\epsilon_\mathrm{p}(\omega)}{|\epsilon_\mathrm{s}-\epsilon_\mathrm{p}(\omega)|^2} > 0. \tag{2.11}$$

Thus $-1/\lambda(\omega)$ is nonreal and $I+\lambda K$ is invertible for every fixed $d>0$ and every frequency with strictly positive loss. This pointwise statement is not a uniform conditioning estimate: the resolvent norm may still become large near a weakly damped resonance or as the geometry degenerates.

Close-inclusion asymptotics provide a useful geometric benchmark. In high-contrast limiting conductivity problems for two disks, the potential admits a decomposition of the form [2, 14, 15]

$$\Phi(\mathbf{x}) = a\,h(\mathbf{x}) + b(\mathbf{x}), \tag{2.12}$$

where h is generated by two logarithmic sources at the reflection fixed points,

$$h(\mathbf{x}) = \frac{1}{2\pi}\left(\ln|\mathbf{x}-\mathbf{p}_1| - \ln|\mathbf{x}-\mathbf{p}_2|\right). \tag{2.13}$$

The focal-point distance from the gap center is $\mathcal{O}(\sqrt{rd})$, and $\|\nabla h\|_{L^\infty}$ scales like $d^{-1/2}$. In that limiting problem one obtains

$$\|\nabla\Phi\|_{L^\infty} = \mathcal{O}\left(|a|d^{-1/2}\right), \qquad d\to 0. \tag{2.14}$$

The cited result is not, by itself, a proof of the same asymptotic law for a fixed lossy Drude–Lorentz material. In the present problem the amplitude and even the effective gap scaling depend on material contrast and loss [22]. We use the limiting result only to motivate local refinement and do not invoke it in the reduced-basis error analysis.

2.3. RCIP revisit. For each fixed $d > 0$, the boundary is smooth and disjoint, so K is compact on the symmetrized space discussed above. The difficulty as d decreases is numerical rather than a loss of compactness at a fixed geometry: interactions between facing boundary segments become sharply localized and increasingly difficult to resolve.

For each physical gap g, let $J_g = J_{g,-} \cup J_{g,+} \subset \Gamma$ be the union of the two facing boundary patches. For configurations with several gaps, set $J^\star = \bigcup_g J_g$ and let $\chi_\star$ be its indicator. We define

$$(K^\star\rho)(\mathbf{x}) = 2\chi_\star(\mathbf{x}) \int_\Gamma \chi_\star(\mathbf{y}) \frac{\partial G(\mathbf{x},\mathbf{y})}{\partial \mathbf{n}_\mathbf{x}} \rho(\mathbf{y})\, ds_\mathbf{y}, \qquad K^\circ = K - K^\star. \tag{2.15}$$

Thus $K^\star$ contains the self- and cross-interactions of both sides of every selected gap. If local neighborhoods overlap, their union is taken before the mask is applied so that entries are not counted twice. This is a discretization-oriented splitting, not a decomposition into noncompact and compact operators.

We then define a transformed density as:

$$\tilde{\rho}(r) = (I + \lambda K^\star)\rho(r). \tag{2.16}$$

Denoting the inverse operator $(I + \lambda K^\star)^{-1}$ by R, the original problem is recast as:

$$(I + \lambda K^\circ R)\, \tilde{\rho}(r) = \lambda g(r), \quad r \in \Gamma. \tag{2.17}$$

The inverse R contains the local fine-scale correction associated with the chosen split. RCIP approximates its action by a local recursion and compresses that action onto the coarse grid. We write $R(\lambda)$ and $\mathbf{R}(\lambda)$ for the continuous notation and its discrete compressed representation, respectively; no spectral improvement theorem for $K^\circ R$ is required below.

2.4. RCIP discretization on coarse and fine grids. The Nyström-based RCIP scheme introduces two underlying hierarchical meshes: a coarse mesh and a fine mesh. The coarse mesh consists of uniform panels discretized via composite Gauss-Legendre quadrature. A strict topological requirement is enforced such that the nearly singular points γ_j coincide exactly with the panel endpoints (junctions). The four panels closest to γ_j (two on each side of the gap) are denoted by $\Gamma^{j\star}$. The fine mesh is subsequently generated from the coarse mesh by applying repeated dyadic subdivisions specifically to the two panels closest to each γ_j. The first row in Figure 5 provides an illustrative example of this mesh architecture.

Let $I^\star_{\text{coa}}$ and $I^\star_{\text{fin}}$ denote the coarse- and fine-grid nodes in these interaction neighborhoods. The fine set contains the nodes introduced by recursive dyadic refinement, whereas the coarse set contains the nodes on the original interaction panels. The complementary index sets are I°_{coa} and I°_{fin}. Outside the refined neighborhoods the two meshes use the same physical nodes. Their total sizes are N_{g} and N_{fin}.

The discrete local index set contains all quadrature nodes in $J^\star$, including the facing patches on both sides of each gap. If $\mathbf{P}_\star$ is the diagonal projector onto this union, the matrix splitting is

$$\mathbf{K}^\star = \mathbf{P}_\star \mathbf{K} \mathbf{P}_\star, \qquad \mathbf{K}^\circ = \mathbf{K} - \mathbf{K}^\star. \tag{2.18}$$

This definition retains the cross-gap entries of $\mathbf{K}^\star$. Several local blocks may be processed independently only when the corresponding off-block interactions have been assigned consistently to $\mathbf{K}^\circ$.

Discretizing (2.7) on the locally refined mesh gives

$$(\mathbf{I}_{\mathrm{fin}} + \lambda \mathbf{K}_{\mathrm{fin}})\, \boldsymbol{\rho}_{\mathrm{fin}} = \lambda \mathbf{g}_{\mathrm{fin}} \tag{2.19}$$

The RCIP compression gives the corresponding coarse-grid system for the transformed density:

$$(\mathbf{I}_{\mathrm{coa}} + \lambda \mathbf{K}^{\circ}_{\mathrm{coa}} \mathbf{R}(\lambda))\, \tilde{\boldsymbol{\rho}}_{\mathrm{coa}}(\lambda) = \lambda \mathbf{g}_{\mathrm{coa}}. \tag{2.20}$$

The purpose of the transformation is to move the fine-scale response generated by $K^\star$ into the compressed inverse. RCIP then represents $\tilde{\rho}$ on the coarse grid while resolving the local inverse on nested fine panels. Piecewise smoothness of the computed transformed density is an observed and intended property of the RCIP construction; the adequacy of a particular coarse grid remains a discretization question and is not implied solely by (2.17).

At the discrete level, $\mathbf{R}(\lambda)$ approximates the action of the fine-grid inverse $(\mathbf{I}_{\mathrm{fin}} + \lambda \mathbf{K}^{\star}_{\mathrm{fin}})^{-1}$ after prolongation and weighted compression. The compression is high order but not algebraically lossless. The local matrices are built by a forward recursion on nested meshes and compressed to the coarse interaction panels. Supp. section SM3 summarizes this construction.

2.5. Empirical motivation for reducing the RCIP state. The original fine-grid equation is affine in λ, whereas the RCIP coarse-grid equation contains the non-affine matrix $\mathbf{R}(\lambda)$. The reason for reducing the latter is empirical rather than spectral: the transformed density is resolved on a much smaller coarse grid and its snapshot matrix is more compressible in the experiments reported here. Figure 2 compares the singular values of the two snapshot matrices, and Figure 3 illustrates the different local variation of the corresponding densities. These figures provide evidence for the tested geometries; they do not establish a general rate for the Kolmogorov N-width.

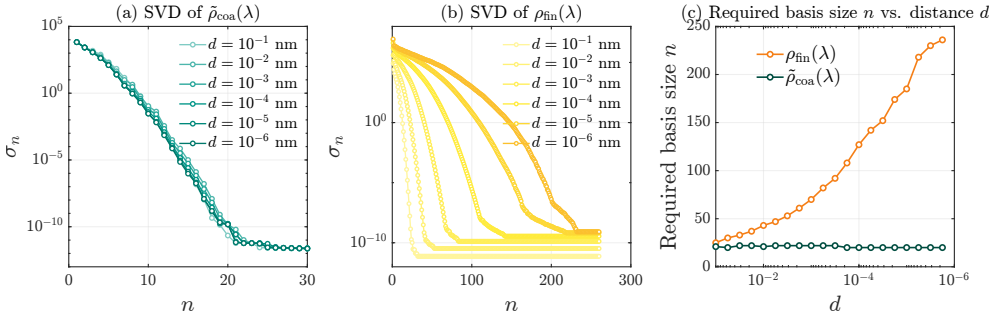

FIG. 2. *Normalized singular values of snapshot matrices for the physical density on the locally refined grid and the RCIP-transformed density on the coarse grid. For the sampled frequencies and geometries shown, the transformed-density snapshots exhibit faster decay. Because the two vectors live on different discretizations, this comparison is a numerical compressibility diagnostic rather than an equality of approximation errors in a common continuous norm.*

3. RCIP-Based Reduced-Basis Approximation.

3.1. Fixed discrete space and Galerkin projection. Throughout this section the geometry, coarse Nyström grid, and RCIP splitting are fixed; $\lambda = \lambda(\omega)$ is the only online parameter. We abbreviate $N = N_{\mathrm{g}}$ and write (2.20) as

$$\mathbf{A}(\lambda)\tilde{\boldsymbol{\rho}}(\lambda) = \mathbf{b}(\lambda), \qquad \mathbf{A}(\lambda) = \mathbf{I} + \lambda \mathbf{K}^{\circ}\mathbf{R}(\lambda), \qquad \mathbf{b}(\lambda) = \lambda \mathbf{g}. \tag{3.1}$$

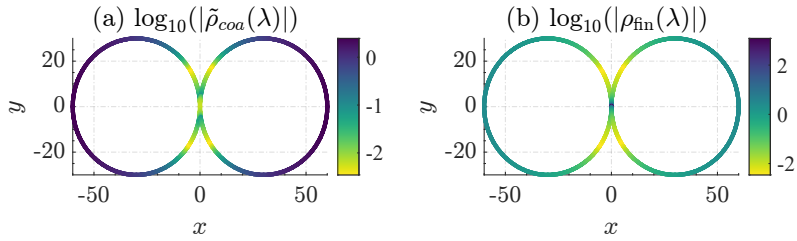

FIG. 3. *Physical fine-grid density* $\boldsymbol{\rho}_{\text{fin}}$ *and transformed coarse-grid density* $\tilde{\boldsymbol{\rho}}_{\text{coa}}$ *for* $d = 10^{-4}$ *nm and photon energy* $\hbar\omega = 3$ *eV. The plot illustrates the localization resolved by the RCIP transformation; the two curves represent different discrete unknowns and are not compared pointwise as approximations of the same function.*

Set $\mathbf{M} = \mathbf{W}_{\text{coa}}$, the positive diagonal coarse-grid quadrature matrix. The parameter-independent inner product and norm are

$$\langle \mathbf{u}, \mathbf{v} \rangle_{\mathbf{M}} = \mathbf{u}^{\mathrm{H}} \mathbf{M} \mathbf{v}, \qquad \|\mathbf{v}\|_{\mathbf{M}} = (\mathbf{v}^{\mathrm{H}} \mathbf{M} \mathbf{v})^{1/2}. \tag{3.2}$$

Because $\mathbf{M}$ is Hermitian positive definite and does not depend on λ, (3.2) is a valid inner product on $\mathbb{C}^N$. In contrast, $\mathbf{u}^{\mathrm{H}} \mathbf{W}_{\text{coa}} \mathbf{R}(\lambda) \mathbf{v}$ is a useful output pairing in some RCIP quadratures but is not, in general, an inner product: $\mathbf{R}(\lambda)$ need not be Hermitian or positive definite.

For snapshots $\mathbf{S} = [\tilde{\boldsymbol{\rho}}(\lambda_1), \ldots, \tilde{\boldsymbol{\rho}}(\lambda_{N_{\text{s}}})]$, a weighted POD basis may be obtained from

$$\mathbf{M}^{1/2} \mathbf{S} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^{\mathrm{H}}, \qquad \boldsymbol{\Phi} = \mathbf{M}^{-1/2} \mathbf{U}(:, 1:n), \qquad \boldsymbol{\Phi}^{\mathrm{H}} \mathbf{M} \boldsymbol{\Phi} = \mathbf{I}_n. \tag{3.3}$$

This basis minimizes the summed $\mathbf{M}$-norm projection error over the sampled snapshots, with minimum $\sum_{j>n} \sigma_j^2$. This finite-snapshot optimality does not imply a decay rate for the Kolmogorov n-width of the continuous solution manifold. In the residual-greedy construction used below, the columns of $\boldsymbol{\Phi}$ instead span selected snapshots and are $\mathbf{M}$-orthonormalized. Snapshot singular values are therefore used as a compressibility diagnostic rather than as a definition of the greedy basis.

We approximate the transformed density by

$$\tilde{\boldsymbol{\rho}}_n(\lambda) = \boldsymbol{\Phi} \mathbf{c}(\lambda) \tag{3.4}$$

and impose the weighted Galerkin condition

$$\boldsymbol{\Phi}^{\mathrm{H}} \mathbf{M} [\mathbf{b}(\lambda) - \mathbf{A}(\lambda) \boldsymbol{\Phi} \mathbf{c}(\lambda)] = \mathbf{0}. \tag{3.5}$$

The reduced system is

$$\mathbf{A}_n(\lambda) \mathbf{c}(\lambda) = \mathbf{b}_n(\lambda), \quad \mathbf{A}_n(\lambda) = \boldsymbol{\Phi}^{\mathrm{H}} \mathbf{M} \mathbf{A}(\lambda) \boldsymbol{\Phi}, \quad \mathbf{b}_n(\lambda) = \lambda \boldsymbol{\Phi}^{\mathrm{H}} \mathbf{M} \mathbf{g}. \tag{3.6}$$

The matrices are complex and generally nonnormal. Invertibility of $\mathbf{A}(\lambda)$ does not imply uniform stability of $\mathbf{A}_n(\lambda)$; the reduced matrix must be nonsingular on the parameter set under consideration, and its conditioning should be monitored in computation.

3.2. Parameter Decoupling of the RCIP Operator. Let $I^\star = \{i_1, \ldots, i_s\}$ be the fixed set of coarse-grid indices covered by the local RCIP blocks, and let $\mathbf{S}_\star \in \{0,1\}^{N \times s}$ be the corresponding extension matrix,

$$(\mathbf{S}_\star)_{i_\ell, \ell} = 1, \qquad \mathbf{S}_\star^{\mathrm{T}} \mathbf{S}_\star = \mathbf{I}_s, \qquad \mathbf{P}_\star = \mathbf{S}_\star \mathbf{S}_\star^{\mathrm{T}}. \tag{3.7}$$

Thus $\mathbf{S}_\star^{\mathrm{T}}$ restricts a global vector to $I^\star$, and $\mathbf{S}_\star$ extends a local vector by zero. The offline–online decomposition uses the following discrete locality property: there exists a parameter-independent matrix $\mathbf{R}_0$ such that

$$\mathbf{R}(\lambda) - \mathbf{R}_0 = \mathbf{P}_\star[\mathbf{R}(\lambda) - \mathbf{R}_0]\mathbf{P}_\star. \tag{3.8}$$

Equivalently,

$$\mathbf{R}(\lambda) = \mathbf{R}_0 + \mathbf{S}_\star \mathbf{D}(\lambda)\mathbf{S}_\star^{\mathrm{T}}, \qquad \mathbf{D}(\lambda) = \mathbf{S}_\star^{\mathrm{T}}[\mathbf{R}(\lambda) - \mathbf{R}_0]\mathbf{S}_\star \in \mathbb{C}^{s\times s}. \tag{3.9}$$

When $\mathbf{R}(\lambda)$ is the identity outside the local blocks, one may set $\mathbf{R}_0 = \mathbf{I}$ and interpret $\mathbf{D}(\lambda)$ as the local correction to the identity. Equation (3.8) is an explicit algebraic assumption, not a consequence of sparsity alone. If it is enforced only approximately, the omitted block must be retained as an additional operator-approximation error.

Define the parameter-independent quantities

$$\mathbf{\Phi}_\star = \mathbf{S}_\star^{\mathrm{T}}\mathbf{\Phi}, \quad \mathbf{H}_0 = \mathbf{\Phi}^{\mathrm{H}}\mathbf{M}\mathbf{K}^\circ\mathbf{R}_0\mathbf{\Phi}, \quad \mathbf{G}_\star = \mathbf{\Phi}^{\mathrm{H}}\mathbf{M}\mathbf{K}^\circ\mathbf{S}_\star, \quad \mathbf{f} = \mathbf{\Phi}^{\mathrm{H}}\mathbf{M}\mathbf{g}. \tag{3.10}$$

Substitution of (3.9) into (3.6) gives

$$\boxed{\mathbf{A}_n(\lambda) = \mathbf{I}_n + \lambda[\mathbf{H}_0 + \mathbf{G}_\star\mathbf{D}(\lambda)\mathbf{\Phi}_\star], \qquad \mathbf{b}_n(\lambda) = \lambda\mathbf{f}.} \tag{3.11}$$

All contractions involving the global grid are parameter independent. For a dense local block, applying $\mathbf{D}(\lambda)$ to $\mathbf{\Phi}_\star$ costs $\mathcal{O}(s^2 n)$ and completing the reduced contraction costs $\mathcal{O}(sn^2)$. These terms cannot in general be reduced to $\mathcal{O}(sn + n^2)$ without additional structure in $\mathbf{D}(\lambda)$.

3.3. Discrete residual indicator and greedy basis construction. For a reduced solution $\mathbf{c}(\lambda)$, define the residual of the full discrete RCIP system by

$$\mathbf{r}_n(\lambda) = \lambda\mathbf{g} - [\mathbf{I} + \lambda\mathbf{K}^\circ\mathbf{R}(\lambda)]\mathbf{\Phi}\mathbf{c}(\lambda). \tag{3.12}$$

Set

$$\mathbf{B}_0 = \mathbf{K}^\circ\mathbf{R}_0\mathbf{\Phi}, \qquad \mathbf{F}_\star = \mathbf{K}^\circ\mathbf{S}_\star, \qquad \mathbf{z}(\lambda) = \mathbf{D}(\lambda)\mathbf{\Phi}_\star\mathbf{c}(\lambda). \tag{3.13}$$

Introduce the parameter-independent residual basis

$$\mathcal{B} = [\,\mathbf{g},\ \mathbf{\Phi},\ \mathbf{B}_0,\ \mathbf{F}_\star\,] \in \mathbb{C}^{N\times p}, \qquad p = 1 + 2n + s. \tag{3.14}$$

Then

$$\mathbf{r}_n(\lambda) = \mathcal{B}\boldsymbol{\zeta}(\lambda), \qquad \boldsymbol{\zeta}(\lambda) = \begin{bmatrix}\lambda & -\mathbf{c}(\lambda)^{\mathrm{T}} & -\lambda\mathbf{c}(\lambda)^{\mathrm{T}} & -\lambda\mathbf{z}(\lambda)^{\mathrm{T}}\end{bmatrix}^{\mathrm{T}}. \tag{3.15}$$

The Hermitian Gram identity is

$$\|\mathbf{r}_n(\lambda)\|_{\mathbf{M}}^2 = \boldsymbol{\zeta}(\lambda)^{\mathrm{H}}(\mathcal{B}^{\mathrm{H}}\mathbf{M}\mathcal{B})\boldsymbol{\zeta}(\lambda). \tag{3.16}$$

Thus every transpose in a squared norm is a conjugate transpose, diagonal terms involving λ contain $|\lambda|^2$, and paired cross terms contain $2\,\mathrm{Re}(\mathbf{v}_a^{\mathrm{H}}\mathbf{M}\mathbf{v}_b)$. Direct evaluation of the squared Gram form can nevertheless lose relative accuracy through cancellation when the residual is small. We therefore compute offline a column-pivoted weighted QR factorization

$$\mathbf{M}^{1/2}\mathcal{B}\mathbf{\Pi} = \mathbf{Q}\mathbf{T}, \qquad \mathbf{Q}^{\mathrm{H}}\mathbf{Q} = \mathbf{I}, \tag{3.17}$$

where $\mathbf{\Pi}$ is a permutation matrix. No column truncation is made in the following identity; a numerical-rank truncation would introduce an additional approximation. It follows that

$$\boxed{\|\mathbf{r}_n(\lambda)\|_{\mathbf{M}} = \|\mathbf{T}\mathbf{\Pi}^{\mathrm{T}}\boldsymbol{\zeta}(\lambda)\|_2.} \tag{3.18}$$

For $\lambda \neq 0$, the relative residual indicator is

$$\eta_{\mathrm{res}}(\lambda) = \frac{\|\mathbf{T}\mathbf{\Pi}^{\mathrm{T}}\boldsymbol{\zeta}(\lambda)\|_2}{|\lambda|(\mathbf{g}^{\mathrm{H}}\mathbf{M}\mathbf{g})^{1/2}}. \tag{3.19}$$

Equations (3.18)–(3.19) evaluate the residual of the discrete RCIP system, up to floating-point error and subject to the locality assumption (3.8).

The residual is not a certified state-error estimator. Indeed, with $\mathbf{e}_n = \tilde{\boldsymbol{\rho}} - \tilde{\boldsymbol{\rho}}_n$,

$$\mathbf{A}(\lambda)\mathbf{e}_n(\lambda) = \mathbf{r}_n(\lambda), \tag{3.20}$$

and, defining

$$\beta_{\mathbf{M}}(\lambda) = \sigma_{\min}(\mathbf{M}^{1/2}\mathbf{A}(\lambda)\mathbf{M}^{-1/2}), \tag{3.21}$$

one obtains

$$\|\mathbf{e}_n(\lambda)\|_{\mathbf{M}} \leq \frac{\|\mathbf{r}_n(\lambda)\|_{\mathbf{M}}}{\beta_{\mathbf{M}}(\lambda)}. \tag{3.22}$$

No computable uniform lower bound for $\beta_{\mathbf{M}}$ is supplied here. Near a weakly damped resonance, a small residual therefore need not imply a proportionally small state or output error.

We use η_{res} to drive basis enrichment on a finite training set Ξ_{train}. The construction is:

1. Select an initial $\lambda_1 \in \Xi_{\mathrm{train}}$, compute one RCIP snapshot, and normalize it in the $\mathbf{M}$-norm.
2. At iteration k, solve (3.11) and evaluate (3.19) for every $\lambda \in \Xi_{\mathrm{train}}$.
3. Choose
$$\lambda_{k+1} \in \arg\max_{\lambda \in \Xi_{\mathrm{train}}} \eta_{\mathrm{res}}(\lambda). \tag{3.23}$$
4. If the maximum indicator is below the prescribed residual tolerance, terminate. Otherwise compute the RCIP snapshot at λ_{k+1}, enrich the space using reorthogonalized $\mathbf{M}$-weighted modified Gram–Schmidt, and update the offline quantities.

The stopping criterion controls only the discrete residual on Ξ_{train}. It does not certify the error between training parameters, the RCIP discretization error, or the error in a physical output. Any tendency of selected frequencies to cluster near resonant features is a numerical observation, not a property implied by the greedy construction.

4. Output Functionals: Efficient Evaluation of Potential and Field Enhancement. The reduced state is useful only if the required layer potentials can also be evaluated without reconstructing a global fine-grid density. We distinguish fixed targets that are far from the boundary, close to a smooth panel, or close to a locally refined gap. Target-dependent factors may be precomputed for prescribed observation points; a new target requires construction of a new quadrature row or direct layer-potential evaluation.

4.1. Product integration. Near-boundary evaluation is difficult because the layer-potential kernels are nearly singular even when the target is not on Γ. We use panelwise product integration following [5, 11]. Let $\mathbf{y}_j : [-1, 1] \to \Lambda_j \subset \Gamma$ be the actual curved-panel parametrization and let $J_j(t) = |\mathbf{y}_j'(t)|$. The scalar single-layer contribution is

$$I_{U,j}(\mathbf{z}) = \int_{-1}^{1} G(\mathbf{z}, \mathbf{y}_j(t))\rho(\mathbf{y}_j(t))J_j(t)\, dt. \tag{4.1}$$

The corresponding field contribution is obtained by replacing G with $\nabla_{\mathbf{z}}G$. This real-arclength representation is used consistently below; no factor $d\tau = i\mathbf{n}\, ds$ is introduced into the physical single-layer density.

For the reported discretization, every panel has $q = 22$ Gauss–Legendre nodes t_a. Let $\mathbf{V}_j \in \mathbb{R}^{q\times q}$, $(\mathbf{V}_j)_{a,k+1} = t_a^k$, and interpolate the density on panel j as

$$\rho(\mathbf{y}_j(t)) \approx \sum_{k=0}^{q-1} \alpha_{j,k} t^k, \qquad \boldsymbol{\alpha}_j = \mathbf{V}_j^{-1}\boldsymbol{\rho}_j. \tag{4.2}$$

For a fixed target, define the curved-panel moment row and field-moment matrix by

$$(\mathbf{m}_{U,j})_{k+1} = \int_{-1}^{1} G(\mathbf{z}, \mathbf{y}_j(t))t^k J_j(t)\, dt, \tag{4.3}$$

$$(\mathbf{M}_{E,j})_{:,k+1} = \int_{-1}^{1} \nabla_{\mathbf{z}}G(\mathbf{z}, \mathbf{y}_j(t))t^k J_j(t)\, dt, \qquad k = 0, \dots, q-1. \tag{4.4}$$

Then

$$I_{U,j}(\mathbf{z}) \approx \mathbf{m}_{U,j}\mathbf{V}_j^{-1}\boldsymbol{\rho}_j, \qquad \mathbf{I}_{E,j}(\mathbf{z}) \approx \mathbf{M}_{E,j}\mathbf{V}_j^{-1}\boldsymbol{\rho}_j. \tag{4.5}$$

For well-separated targets these weights reduce to ordinary Gauss–Legendre quadrature. For close targets they are evaluated with the singularity-subtraction and branch-consistent recurrences of [5, 11]. The Jacobian J_j and the curved map $\mathbf{y}_j$ remain in the pulled-back integral; replacing them by the endpoint chord would be valid only for a straight panel. Supp. section SM2 records the canonical Cauchy and logarithmic moments used as building blocks.

4.2. Field evaluation for different targets. For a fixed target $\mathbf{z}$, let $\mathbf{F}_U(\mathbf{z})$ and $\mathbf{F}_E(\mathbf{z})$ denote the discrete potential and field layer-potential operators. The latter has two rows unless the two field components are encoded by one complex number. With $\tilde{\boldsymbol{\rho}}_n = \boldsymbol{\Phi}\mathbf{c}(\lambda)$, the total reduced potential and field are

$$\Phi_n(\mathbf{z}; \lambda) = \Phi_{\text{in}}(\mathbf{z}) + \mathbf{F}_U(\mathbf{z})\mathbf{R}(\lambda)\boldsymbol{\Phi}\mathbf{c}(\lambda), \tag{4.6}$$

$$\mathbf{E}_n(\mathbf{z}; \lambda) = \mathbf{E}_0 - \mathbf{F}_E(\mathbf{z})\mathbf{R}(\lambda)\boldsymbol{\Phi}\mathbf{c}(\lambda). \tag{4.7}$$

The field enhancement $\eta_n = \|\mathbf{E}_n\|_2/\|\mathbf{E}_0\|_2$ is a nonlinear postprocessing quantity, not a linear functional of the density. Supp. section SM6 derives the discrete duality pairing that contains $\mathbf{R}(\lambda)$. The quadrature treatment depends on the target location (see Figure 4):

1. **Type 1 and 2 targets**: Standard coarse-grid quadrature is used away from the boundary, and product integration is used near a smooth panel away from a gap.
2. **Type 3 targets**: A local backward RCIP recursion reconstructs the physical density on the refined gap panels before product integration; see Supp. section SM5. Only the local part is reconstructed.

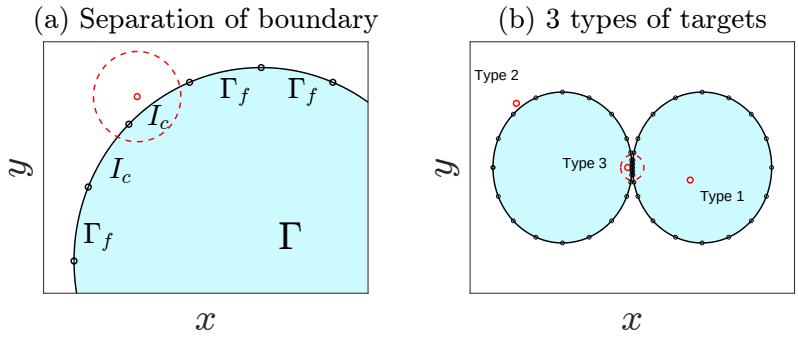

FIG. 4. *Separation of boundary for product integration and three different types of targets that require distinct numerical treatments during integration.*

4.3. Potential for targets of Type 1 and 2. The layer-potential contribution at a fixed Type 1 or 2 target has the form

$$F_U^{(i)}(\tilde{\boldsymbol{\rho}}_{\mathrm{coa}};\lambda) = \mathbf{F}_{U,\mathrm{coa}}^{(i)}\mathbf{R}(\lambda)\tilde{\boldsymbol{\rho}}_{\mathrm{coa}} = \mathbf{F}_{U,\mathrm{coa}}^{(i)}\mathbf{R}(\lambda)\boldsymbol{\Phi}\mathbf{c}_\lambda, \quad i=1,2. \tag{4.8}$$

where $\mathbf{F}_{U,\mathrm{coa}}^{(i)} \in \mathbb{C}^{1\times N}$ is the target-dependent quadrature row. In the panelwise product-integration notation,

$$\mathbf{F}_{U,\mathrm{coa}}^{(i)} = \left[\mathbf{m}_{U,1}^{(i)}\mathbf{V}_1^{-1} \ \cdots \ \mathbf{m}_{U,N_p}^{(i)}\mathbf{V}_{N_p}^{-1}\right], \qquad N_p = N/q. \tag{4.9}$$

Here $\mathbf{m}_{U,j}^{(i)}$ is the curved-panel moment row in (4.4) for target i. This block assembly avoids interpreting one $q\times q$ Vandermonde matrix as a global operator.

Using (3.9), the reduced layer-potential contribution is

$$\mathbf{F}_{U,\mathrm{coa}}^{(i)}\mathbf{R}(\lambda)\boldsymbol{\Phi}\mathbf{c}(\lambda) = \underbrace{\mathbf{F}_{U,\mathrm{coa}}^{(i)}\mathbf{R}_0\boldsymbol{\Phi}}_{\text{offline}}\mathbf{c}(\lambda) + \underbrace{\mathbf{F}_{U,\mathrm{coa}}^{(i)}\mathbf{S}_\star}_{\text{offline}}\mathbf{D}(\lambda)\boldsymbol{\Phi}_\star\mathbf{c}(\lambda). \tag{4.10}$$

For a prescribed target, all global-grid factors in (4.10) are precomputed. The total potential is obtained by adding Φ_{in} as in (4.7). Changing the target changes $\mathbf{F}_{U,\mathrm{coa}}^{(i)}$ and therefore requires new target-specific setup.

4.4. Field enhancement for targets of Type 3. For a Type 3 target, the field layer-potential term is evaluated from the locally reconstructed physical density:

$$\mathbf{F}_E^{(3)}(\mathbf{z})\boldsymbol{\rho}_{\mathrm{fin},n}(\lambda). \tag{4.11}$$

The target-dependent field integration matrix on the refined mesh is

$$\mathbf{F}_{E,\mathrm{fin}}^{(3)} = \left[\mathbf{M}_{E,1}^{(3)}\mathbf{V}_1^{-1} \ \cdots \ \mathbf{M}_{E,N_p^{\mathrm{fin}}}^{(3)}\mathbf{V}_{N_p^{\mathrm{fin}}}^{-1}\right] \in \mathbb{C}^{2\times N_{\mathrm{fin}}}, \qquad N_p^{\mathrm{fin}} = N_{\mathrm{fin}}/q. \tag{4.12}$$

Each $\mathbf{M}_{E,j}^{(3)} \in \mathbb{C}^{2\times q}$ is defined by (4.4); its two rows correspond to the Cartesian components of $\nabla_{\mathbf{z}}G$. The same definition applies if the implementation stores these components in an equivalent complex encoding.

Let $\mathbf{S}_{\circ,\mathrm{coa}}$ and $\mathbf{S}_{\circ,\mathrm{fin}}$ inject the coincident nonlocal nodes on the coarse and fine meshes, and let $\mathbf{S}_{\star,\mathrm{coa}}$ and $\mathbf{S}_{\star,\mathrm{fin}}$ inject the corresponding local nodes. The backward recursion in Supp. section SM5 defines a local map $\mathcal{P}_{\lambda,\star}$, giving

$$\boldsymbol{\rho}_{\mathrm{fin},n}(\lambda) = \mathbf{S}_{\circ,\mathrm{fin}}\mathbf{S}_{\circ,\mathrm{coa}}^{\mathrm{T}}\boldsymbol{\Phi}\mathbf{c}(\lambda) + \mathbf{S}_{\star,\mathrm{fin}}\mathcal{P}_{\lambda,\star}\mathbf{S}_{\star,\mathrm{coa}}^{\mathrm{T}}\boldsymbol{\Phi}\mathbf{c}(\lambda). \tag{4.13}$$

The first term in (4.13) is parameter independent apart from $\mathbf{c}(\lambda)$ and can be contracted with $\mathbf{F}_{E,\mathrm{fin}}^{(3)}$ offline. The second term requires a parameter-dependent local backward recursion. Its cost depends on the local refinement depth and local block size, even though it does not involve the global coarse-grid dimension. The incident field is added before forming the nonlinear enhancement in (4.7).

4.5. Operation counts. Let $C_R(s, n_{\mathrm{sub}}; \lambda)$ denote the work required to construct the local correction $\mathbf{D}(\lambda)$ by the forward RCIP recursion, and let $C_P(s, n_{\mathrm{sub}}; \lambda)$ denote the local backward-reconstruction work. These quantities remain part of every query. In an implementation that factors dense blocks of size proportional to s at each level, C_R and C_P typically contain terms proportional to $n_{\mathrm{sub}}s^3$, but their precise constants depend on the local panel topology.

For a fixed final basis and a dense representation of $\mathbf{K}^{\circ}$, global preprocessing includes $\mathcal{O}(N^2n)$ work to form $\mathbf{K}^{\circ}\mathbf{R}_0\boldsymbol{\Phi}$ and up to $\mathcal{O}(Np^2)$ work for the weighted residual factorization, where $p = 1 + 2n + s$. Each selected snapshot additionally requires a high-fidelity RCIP solve. These costs may be reduced by a fast matrix–vector product, but no such acceleration is assumed here.

For one new frequency, reduced assembly and solution cost

$$C_R(s, n_{\mathrm{sub}}; \lambda) + \mathcal{O}(s^2n + sn^2 + n^3). \tag{4.14}$$

Evaluating the QR residual indicator adds at most $\mathcal{O}(p^2)$ work. A fixed Type 1 or 2 output requires only local and reduced contractions once its target row has been constructed. A Type 3 output additionally incurs C_P and integration over the locally refined nodes. Consequently, for fixed s and fixed recursion settings, the post-offline algebra contains no operations on the global coarse grid. It is not independent of n_{sub}, and a new geometry or target requires new offline quantities. The principal per-query terms are summarized in Table 1. The counts describe algebraic work only; the present manuscript does not report measured wall-clock timings.

TABLE 1

Leading algebraic costs per frequency for dense full-order operators. Here k_{fin} and k_{g} are iteration counts for the corresponding full systems, $N_{\mathrm{fin}}^{\star}$ is the number of locally refined nodes used for a Type 3 output, and target rows are assumed fixed.

Method	Local RCIP setup	Linear solve/assembly	Type 1–2 output	Type 3 output
Locally refined system	–	$\mathcal{O}(k_{\mathrm{fin}}N_{\mathrm{fin}}^2)$	$\mathcal{O}(N_{\mathrm{fin}})$	$\mathcal{O}(N_{\mathrm{fin}})$
Standard RCIP	C_R	$\mathcal{O}(k_{\mathrm{g}}N^2)$	$\mathcal{O}(N)$	$C_P + \mathcal{O}(N + N_{\mathrm{fin}}^{\star})$
RCIP reduced basis	C_R	$\mathcal{O}(s^2n + sn^2 + n^3)$	$\mathcal{O}(s^2 + sn + n)$	$C_P + \mathcal{O}(N_{\mathrm{fin}}^{\star} + sn + n)$

The N^2 terms in Table 1 are dense matrix–vector products per iteration. Fast summation could change those terms, so the table should not be read as a timing comparison with optimized full-order solvers.

5. Multi-layer RCIP (mRCIP) for Field Evaluation at Extreme Proximity. Standard RCIP uses a narrow self-similar refinement hierarchy around each difficult boundary junction [10]. In the calculations at the gap center, the reported

output discrepancy deteriorates for the smallest gaps when this standard hierarchy is used; see Figure 8. The observation is consistent with insufficient resolution in the local backward reconstruction and product integration, although the present experiments do not separate all sources of high-fidelity discretization error.

We therefore test a wider local hierarchy, termed multi-layer RCIP (mRCIP). A layer-power parameter l_p controls the lateral span of panels refined at each recursion level. This changes the local discretization and reconstruction only; it is not accompanied here by an a priori error estimate.

In the standard construction used here, n_{sub} controls the recursion depth and $l_p = 2$. Increasing l_p widens the set of panels included in each local level; the topology for $l_p = 3$ is shown in Figure 5. The supplementary material gives a selection-matrix description that allows the panel and node ordering to depend on the chosen local topology.

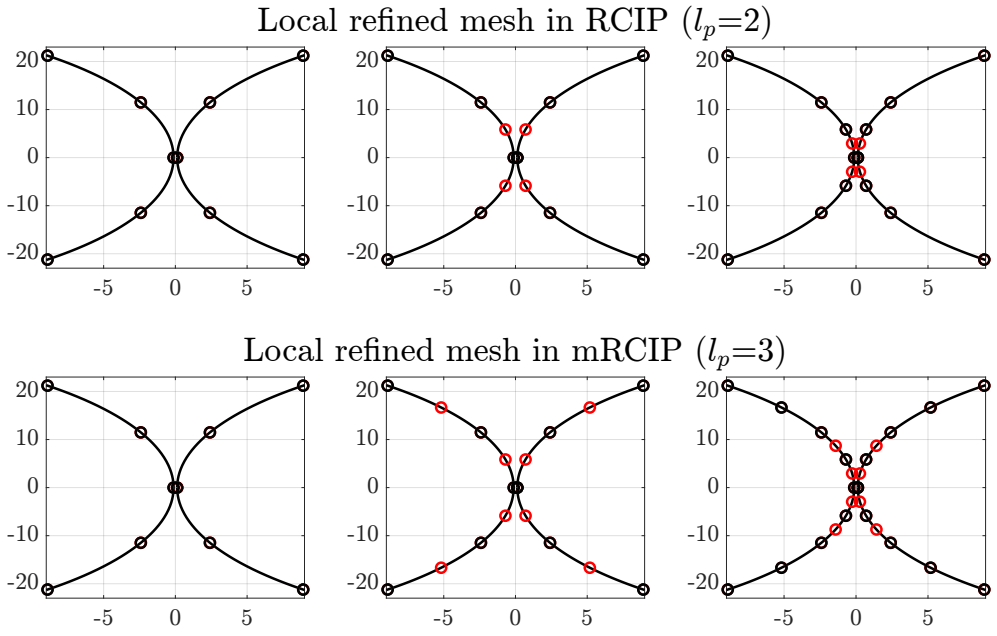

FIG. 5. *Local hierarchical meshes used by standard RCIP ($l_p = 2$) and the wider mRCIP variant. Increasing l_p enlarges the group of panels refined at each level and hence the local block size.*

The selection-matrix formulation in subsection 3.2 also applies to the wider hierarchy after enlarging $I^\star$ and s. The global-grid contractions remain offline, but C_R, C_P, and the local reduced contractions all increase with the block size. Figure 8 reports the observed difference between $l_p = 2$ and $l_p = 3$ at the gap center; the value of l_p is stated for subsequent Type 3 calculations.

6. Numerical Results. The baseline test uses two silver disks of radius $r = 30$ nm and photon energies $\hbar\omega \in [0.1, 4.5]$ eV. No wall-clock measurements are reported, so this section assesses approximation behavior rather than measured speed.

Implementation-status note.. The fixed-$\mathbf{M}$ projection, explicit locality identity, and Hermitian QR residual in section 3 correct formulas used in an earlier version of the manuscript. The existing numerical files must therefore be regenerated, or

independently shown to be unchanged, using the revised implementation before submission. Until that consistency check is complete, the plots below document the scope and outcomes of the legacy calculations but do not validate every algebraic step of the revised formulation.

Scope of the comparisons.. Let $\boldsymbol{u}_h$ denote the solution produced by the designated RCIP discretization and $\boldsymbol{u}_n$ its reduced approximation. Unless stated otherwise, a reported error is a discrete ROM-to-RCIP discrepancy for the same geometry, sampled frequency, and target. It isolates reduction and output-evaluation differences but does not estimate the error of the RCIP solution relative to the continuous boundary integral equation. Where two RCIP reconstruction strategies are compared, the plotted discrepancy is relative to the reference calculation used for that comparison.

A separate basis and separate offline operators are constructed for each gap distance and disk configuration. The gap, radii, particle positions, and prescribed observation targets are therefore not online parameters. The residual-greedy search is evaluated on a finite set of 400 frequencies and supplies no guarantee between those samples. Finally, the smallest gaps are numerical stress tests of the two-dimensional local-dielectric model, not physically complete subnanometer predictions.

6.1. Observed basis-size dependence and Type 1–2 outputs. Figure 6 shows the transformed-density and potential discrepancies as the reduced dimension increases for one fixed geometry and two prescribed targets. The curves decrease rapidly over the displayed range and reach approximately 10^{-10} at $n = 20$ in this test. This is numerical evidence of snapshot compressibility for the sampled problem; it is not a proof of exponential convergence or of a Kolmogorov-width decay rate.

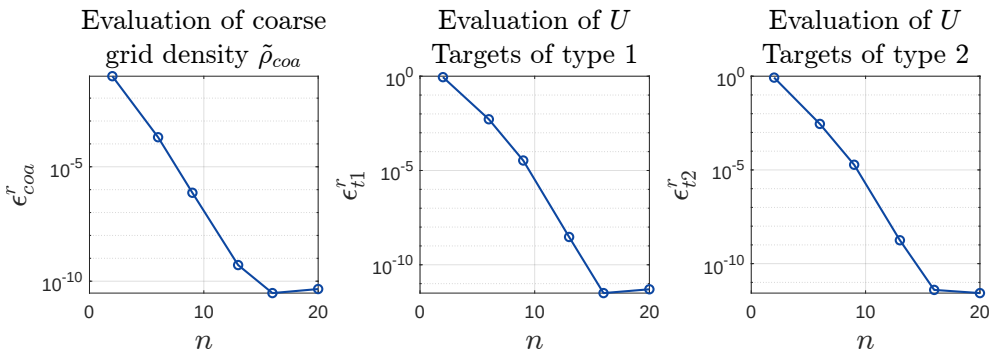

FIG. 6. *Basis-size dependence for $d = 10^{-4}$ nm and $r = 30$ nm. Targets are $25-5i$ (Type 1) and $-51.5085 + 21.5179i$ (Type 2). The plotted quantities are relative Euclidean discrepancies between the reduced and designated RCIP results: $\epsilon^{\mathrm{r}}_{\mathrm{t}i} = \|U_n - U_h\|_2/\|U_h\|_2$ and $\epsilon^{\mathrm{r}}_{\mathrm{coa}} = \|\tilde{\boldsymbol{\rho}}_n - \tilde{\boldsymbol{\rho}}_h\|_2/\|\tilde{\boldsymbol{\rho}}_h\|_2$.*

6.2. Observed effect of multi-layer RCIP. Figure 7 compares the pointwise potential discrepancies produced by the standard hierarchy ($l_p = 2$) and the wider hierarchy ($l_p = 5$) on the same 500×500 target grid. In this test, the wider hierarchy substantially reduces the localized discrepancy near the closest boundary points. The circular transition at radius 10 nm is the prescribed switch to the Type 3 evaluation rule and is therefore an algorithmic feature, not a physical interface. Because the comparison is against a discrete reference calculation, it does not establish an absolute error bound for mRCIP.

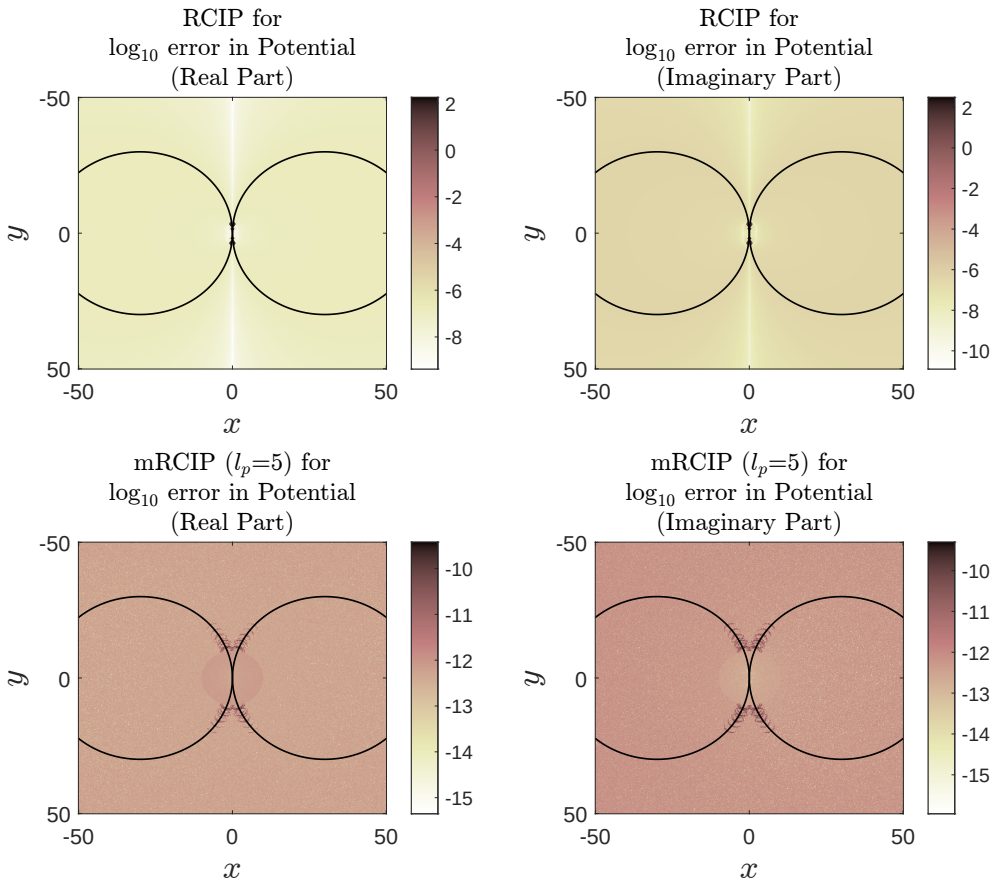

FIG. 7. *Pointwise absolute potential discrepancy at photon energy $\hbar\omega = 2.4$ eV for $r = 30$ nm and $d = 10^{-4}$ nm, relative to the reference calculation used in this comparison. Top: standard RCIP ($l_p = 2$). Bottom: mRCIP ($l_p = 5$). Each disk has 80 coarse panels with $q = 22$ Gauss–Legendre nodes per panel, and the target grid is 500×500. The circle of radius 10 nm marks the prescribed switch to Type 3 evaluation.*

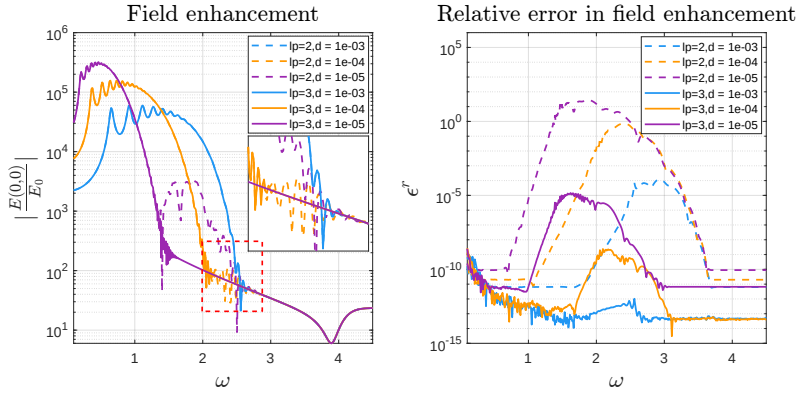

FIG. 8. *Field enhancement at the gap center for standard RCIP ($l_p = 2$) and mRCIP ($l_p = 3$), evaluated at the plotted frequency samples. The curves show the observed sensitivity of the Type 3 output to the local reconstruction width for small gaps.*

6.3. Residual-greedy behavior on the training set. Figure 9 reports the residual indicator, selected snapshot frequencies, and transformed-state discrepancies on the 400-point training set.

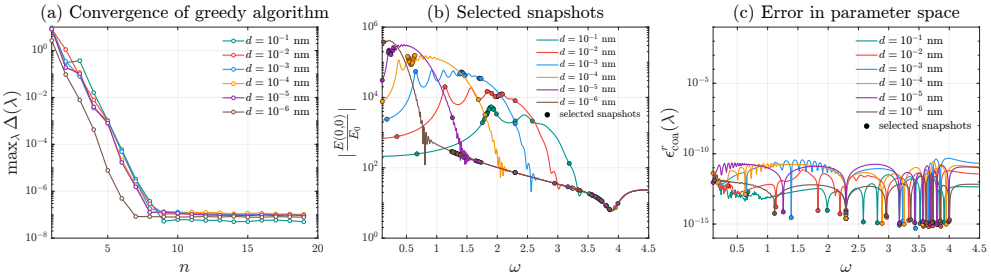

FIG. 9. *Working-draft residual-greedy results on the finite training set. Panel (a) was generated with the legacy residual expression and must be regenerated using* (3.18) *and* (3.19)*; it is not evidence for the revised indicator. Panels (b)–(c) show the legacy selected frequencies and corresponding ROM-to-RCIP transformed-state discrepancies and must likewise be checked after the implementation is reconciled with section* 3.

In the legacy panel (a), the plotted quantity decreases by several orders of magnitude and then plateaus near 10^{-7}–10^{-8} for the displayed cases. Because that panel did not use (3.18), no quantitative conclusion about the revised residual indicator is drawn from the plateau.

The legacy panel (b) places several selected frequencies near prominent resonant features, and panel (c) reports small state discrepancies at the sampled frequencies. These observations may change when the fixed-$\mathbf{M}$ projection and corrected residual are used. Even after regeneration, η_{res} is not divided by a verified stability lower bound, so agreement between residual and state-error trends would remain empirical rather than certified.

6.4. Frequency-sweep comparisons for fixed dimer geometries. For each gap distance, the residual-greedy construction uses 400 frequency samples and stops when the residual tolerance is reached or when $n = 20$. In the displayed tests the dimension cap is reached. The models for different gaps are trained separately.

Figure 10 compares the reduced and RCIP field enhancements at the sampled frequencies for gaps from 10^{-1} to 10^{-6} nm. The reduced curves resolve the displayed resonant features, and the plotted relative discrepancies are below 10^{-5} at most samples. This observation does not establish a uniform error over the interval. The $d = 10^{-6}$ nm case is included only as a stress test of the local-dielectric discretization and reduction; tunneling and nonlocal effects are outside the model [7, 24].

6.5. A fixed asymmetric-dimer test case. We next consider two silver disks with radii $r_1 = 30$ nm and $r_2 = 20$ nm. A separate reduced model is constructed for every gap shown in Figure 12. The reduced curves follow the RCIP reference values at the plotted frequency samples; this test examines one family of asymmetric dimers and does not imply geometry-parametric accuracy.

The spatial potential for $\hbar\omega = 2.5$ eV and $d = 10^{-4}$ nm is shown in Figure 13. On the displayed target grid, the reported pointwise ROM-to-RCIP absolute discrepancy is below 10^{-8}. This is evidence for the stated fixed geometry only.

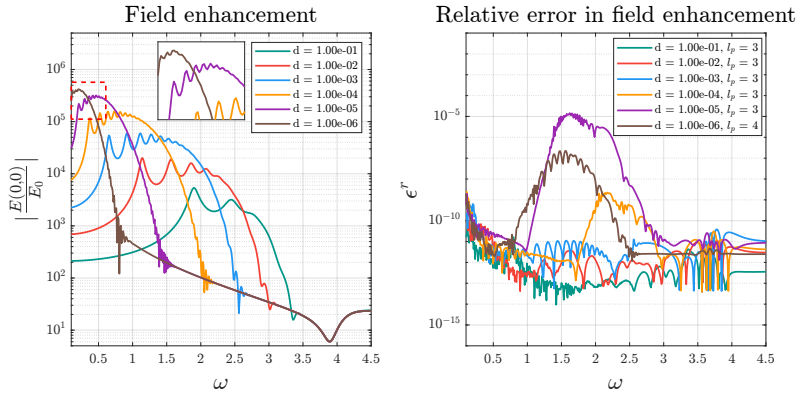

FIG. 10. *Field enhancement at* $(0,0)$ *and ROM-to-RCIP relative discrepancy for separately trained fixed-gap models with* $n = 20$. *The layer power* l_p *is chosen separately for the displayed gap cases; no automatic selection rule or a priori output-error bound is used.*

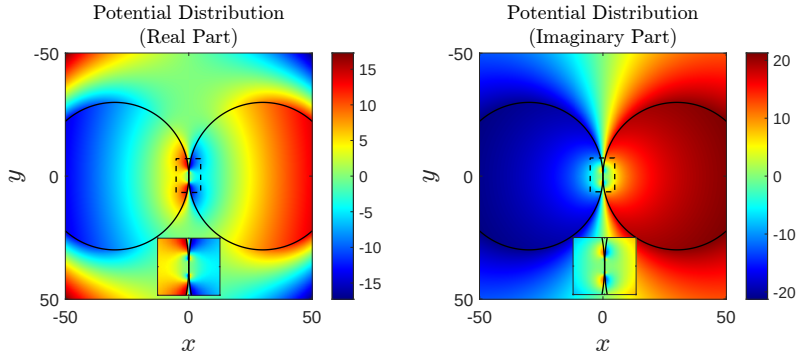

FIG. 11. *Real and imaginary parts of the potential for two silver disks of radius* 30 *nm with* $d = 10^{-4}$ *nm,* $\mathbf{E}_0 = (1,0)$, *and* $\hbar\omega = 2.4$ *eV. The designated RCIP calculation uses 80 panels per disk,* $q = 22$ *Gauss–Legendre nodes per panel, and* $l_p = 5$. *The plotted ROM-to-RCIP pointwise relative discrepancy is below* 10^{-9} *on the target grid.*

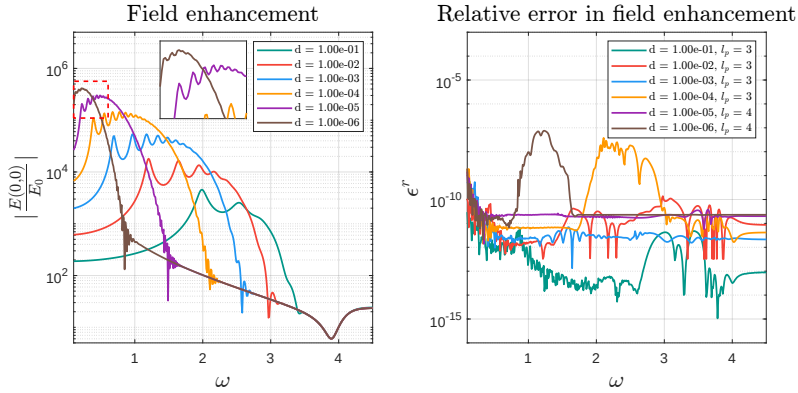

FIG. 12. *Field enhancement at* $(0,0)$ *for separately trained asymmetric-dimer models with* $r_1 = 30$ *nm and* $r_2 = 20$ *nm. The case* $d = 10^{-5}$ *nm uses* $l_p = 4$.

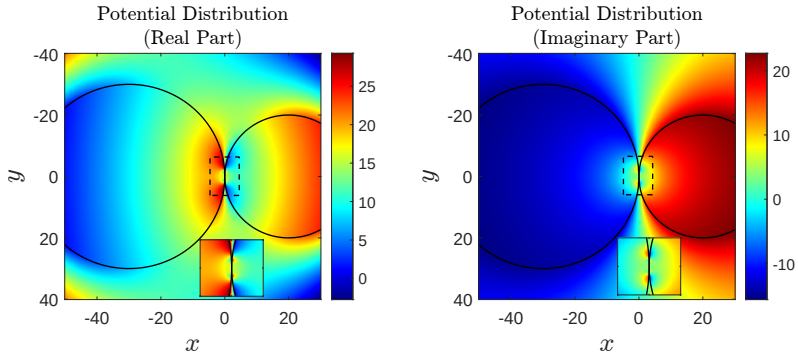

FIG. 13. *Real and imaginary parts of the potential for two asymmetric silver disks ($r_1 = 30$ nm, $r_2 = 20$ nm) with $d = 10^{-4}$ nm, $\mathbf{E}_0 = (1, 0)$, and $\hbar\omega = 2.5$ eV. The plotted pointwise absolute discrepancy is relative to the designated RCIP result.*

6.6. Fixed three-disk test cases. We finally apply the fixed-geometry construction to two configurations of three silver disks with $r = 30$ nm and nominal gap $d = 0.25$ nm.

The local index set is the union of the nodes associated with all close gaps. The local correction $\mathbf{D}(\lambda)$ need not be block diagonal when gap neighborhoods or their interactions couple; the selection-matrix derivation requires only (3.8). Its cost grows with the combined local dimension s, so two trimer examples do not establish asymptotic many-particle scalability.

The first configuration places the disks at the vertices of an equilateral triangle and evaluates the field at the origin. The second rotates the third disk by 80° and is trained independently.

For the symmetric configuration, the reduced and RCIP curves are visually indistinguishable at the plotted samples for $n = 16$, and the reported relative discrepancy is below 10^{-11} there; see Figure 14.

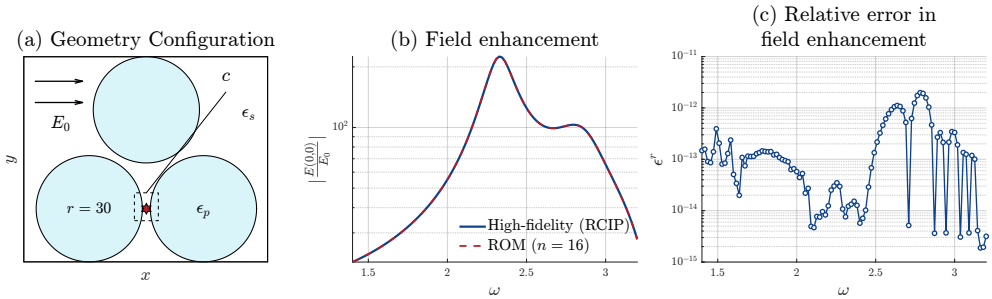

FIG. 14. *Symmetric three-disk test. (a) Three identical disks form an equilateral triangle with $r = 30$ nm and $d = 0.25$ nm. (b) RCIP and reduced-basis field enhancement at the origin for the sampled frequencies. (c) Corresponding ROM-to-RCIP relative discrepancy for $n = 16$.*

For the rotated configuration, a separately trained basis of dimension $n = 20$ gives a reported ROM-to-RCIP field-enhancement discrepancy of order 10^{-10} at the plotted samples; see Figure 15. The altered spectrum is consistent with a different collective response, but a modal interpretation is not developed here.

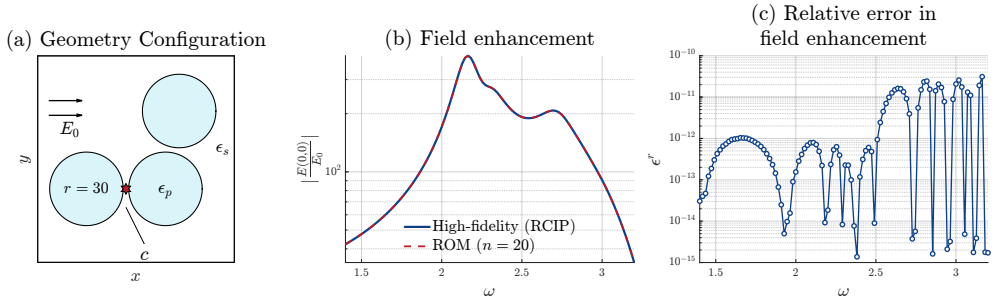

FIG. 15. *Asymmetric three-disk test. (a) The third disk is rotated by* 80°. *(b) RCIP and reduced-basis field enhancement at the origin for the sampled frequencies. (c) Corresponding ROM-to-RCIP relative discrepancy for* $n = 20$.

7. Conclusion and Outlook. We have developed an RCIP-based reduced-basis formulation for repeated frequency-domain solutions of a two-dimensional quasistatic transmission problem with nearly touching disks. The reduced state is the transformed coarse-grid density. For each fixed geometry, the frequency-dependent part of the compressed inverse is represented through local gap indices, which separates global parameter-independent contractions from a local RCIP recursion. The same structure supports weighted Galerkin assembly, a complex discrete residual indicator, and target-specific layer-potential evaluation. A wider local hierarchy is used for prescribed targets close to a gap.

The reported calculations show rapid decay of transformed-density snapshot singular values and small ROM-to-RCIP discrepancies on the sampled frequency sets. Similar behavior is observed for the tested symmetric dimer, asymmetric dimer, and two three-disk configurations. These examples support the use of local RCIP structure in reduced frequency sweeps. They do not establish uniform accuracy on a continuous frequency interval, geometry-independent approximation properties, or asymptotic scalability with particle count.

Several limitations define the present scope. First, the residual indicator is not a certified error estimator because no computable lower bound for $\beta_{\mathbf{M}}(\lambda)$ has been derived near resonance. Second, the reported discrepancies are relative to designated RCIP calculations; the full-order discretization error is not assessed separately. Third, the existing numerical files predate the fixed-$\mathbf{M}$ and Hermitian-QR reformulation and must be reconciled with it before submission. Fourth, each geometry, gap distance, and target setup requires its own reduced model or offline output construction, and the local mRCIP width is selected empirically. Finally, the analysis and experiments use a two-dimensional classical quasistatic model, and the smallest gaps serve only as numerical stress tests.

Natural next steps are an independent RCIP convergence and timing study, verification of the local selection-matrix identity and QR residual against directly assembled vectors, estimation of parameter-dependent stability constants, systematic analysis of l_p, and extension to geometry-parametric and three-dimensional formulations.

Acknowledgments.

REFERENCES

[1] H. Ammari, G. Ciraolo, H. Kang, H. Lee, and G. W. Milton, *Spectral theory of a Neumann-Poincaré-type operator and analysis of cloaking due to anomalous localized resonance*, Arch. Ration. Mech. Anal., 208 (2013), pp. 667–692, https://doi.org/10.1007/s00205-012-0605-5.

[2] H. Ammari, H. Kang, H. Lee, J. Lee, and M. Lim, *Optimal estimates for the electric field in two dimensions*, J. Math. Pures Appl., 88 (2007), pp. 307–324, https://doi.org/10.1016/j.matpur.2007.07.005.

[3] H. Ammari, P. Millien, M. Ruiz, and H. Zhang, *Mathematical analysis of plasmonic nanoparticles: the scalar case*, Arch. Ration. Mech. Anal., 224 (2017), pp. 597–658, https://doi.org/10.1007/s00205-017-1084-5.

[4] H. Ammari, M. Ruiz, S. Yu, and H. Zhang, *Mathematical analysis of plasmonic resonances for nanoparticles: the full Maxwell equations*, J. Differ. Equ., 261 (2016), pp. 3615–3669, https://doi.org/10.1016/j.jde.2016.05.036.

[5] E. Andersson, *Implementation and study of boundary integral operators related to PDEs in the plane*, master's thesis, Lund University, 2023.

[6] J. N. Anker, W. P. Hall, O. Lyandres, N. C. Shah, J. Zhao, and R. P. Van Duyne, *Biosensing with plasmonic nanosensors*, Nat. Mater., 7 (2008), pp. 442–453, https://doi.org/10.1038/nmat2162.

[7] C. Ciracì, R. T. Hill, J. J. Mock, Y. Urzhumov, A. I. Fernández-Domínguez, S. A. Maier, J. B. Pendry, A. Chilkoti, and D. R. Smith, *Probing the ultimate limits of plasmonic enhancement*, Science, 337 (2012), pp. 1072–1074, https://doi.org/10.1126/science.1224823.

[8] M. Fares, J. S. Hesthaven, Y. Maday, and B. Stamm, *The reduced basis method for the electric field integral equation*, J. Comput. Phys., 230 (2011), pp. 5532–5555, https://doi.org/10.1016/j.jcp.2011.03.023.

[9] J. Helsing, *Solving integral equations on piecewise smooth boundaries using the RCIP method: A tutorial*, Abstr. Appl. Anal., 2013 (2013), pp. 1–20, https://doi.org/10.1155/2013/938167.

[10] J. Helsing and A. Karlsson, *On a Helmholtz transmission problem in planar domains with corners*, J. Comput. Phys., 371 (2018), pp. 315–332, https://doi.org/10.1016/j.jcp.2018.05.044.

[11] J. Helsing and R. Ojala, *On the evaluation of layer potentials close to their sources*, J. Comput. Phys., 227 (2008), pp. 2899–2921, https://doi.org/10.1016/j.jcp.2007.11.024.

[12] J. S. Hesthaven, G. Rozza, and B. Stamm, *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*, Springer International Publishing, Cham, 2016, https://doi.org/10.1007/978-3-319-22470-1.

[13] J. S. Hesthaven, B. Stamm, and S. Zhang, *Certified reduced basis method for the electric field integral equation*, SIAM J. Sci. Comput., 34 (2012), pp. A1777–A1799, https://doi.org/10.1137/110848268.

[14] H. Kang, M. Lim, and K. Yun, *Asymptotics and computation of the solution to the conductivity equation in the presence of adjacent inclusions with extreme conductivities*, J. Math. Pures Appl., 99 (2013), pp. 234–249, https://doi.org/10.1016/j.matpur.2012.06.013.

[15] H. Kang, M. Lim, and K. Yun, *Characterization of the electric field concentration between two adjacent spherical perfect conductors*, SIAM J. Appl. Math., 74 (2014), pp. 125–146, https://doi.org/10.1137/130922434.

[16] A. Monje-Real and V. de la Rubia, *Electric field integral equation fast frequency sweep for scattering of nonpenetrable objects via the reduced-basis method*, IEEE Trans. Antennas Propag., 68 (2020), pp. 6232–6244, https://doi.org/10.1109/tap.2020.2992882.

[17] S. Nie and S. R. Emory, *Probing single molecules and single nanoparticles by surface-enhanced Raman scattering*, Science, 275 (1997), pp. 1102–1106, https://doi.org/10.1126/science.275.5303.1102.

[18] J. B. Pendry, A. I. Fernández-Domínguez, Y. Luo, and R. Zhao, *Capturing photons with transformation optics*, Nat. Phys., 9 (2013), pp. 518–522, https://doi.org/10.1038/nphys2667.

[19] A. Quarteroni, A. Manzoni, and F. Negri, *Reduced Basis Methods for Partial Differential Equations*, Springer International Publishing, Cham, 2016, https://doi.org/10.1007/978-3-319-15431-2.

[20] I. Romero, J. Aizpurua, G. W. Bryant, and F. J. García de Abajo, *Plasmons in nearly touching metallic nanoparticles: singular response in the limit of touching dimers*, Opt. Express, 14 (2006), p. 9988, https://doi.org/10.1364/oe.14.009988.

[21] G. Rozza, D. B. P. Huynh, and A. T. Patera, *Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations: Application to transport and continuum mechanics*, Arch. Comput. Methods Eng.,

612 15 (2008), pp. 229–275, https://doi.org/10.1007/s11831-008-9019-9.
613 [22] O. Schnitzer, *Singular perturbations approach to localized surface-plasmon resonance: Nearly*
614 *touching metal nanospheres*, Phys. Rev. B, 92 (2015), https://doi.org/10.1103/physrevb.
615 92.235428.
616 [23] R. Zhao, Y. Luo, A. I. Fernández-Domínguez, and J. B. Pendry, *Description of van*
617 *der Waals interactions using transformation optics*, Phys. Rev. Lett., 111 (2013), https:
618 //doi.org/10.1103/physrevlett.111.033602.
619 [24] J. Zuloaga, E. Prodan, and P. Nordlander, *Quantum description of the plasmon reso-*
620 *nances of a nanoparticle dimer*, Nano Lett., 9 (2009), pp. 887–891, https://doi.org/10.
621 1021/nl803811g.