

Article

Not peer-reviewed version

Symmetry-Resolved Phase Transitions of Electromagnetic Degrees of Freedom Under RIS Control

[Carlos Bousoño-Calzón](#) *

Posted Date: 3 March 2026

doi: 10.20944/preprints202603.0272.v1

Keywords: symmetry-resolved analysis; phase transitions; radiation operators; electromagnetic degrees of freedom; representation theory of compact groups; reconfigurable intelligent surfaces (RIS); mode resolvability



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Article

Symmetry-Resolved Phase Transitions of Electromagnetic Degrees of Freedom Under RIS Control

Carlos Bousoño-Calzón 

Signal Theory and Communications. Univ. Carlos III de Madrid; cbousono@ing.uc3m.es

Abstract

The theory of physical degrees of freedom (DoF) developed by Franceschetti–Migliore–Minero (FMM) establishes a fundamental phase transition in the singular-value spectrum of electromagnetic radiation operators under maximal rotational symmetry. In this work, we revisit this result from a symmetry-explicit operator-theoretic perspective and extend it to scenarios with reduced and controllable symmetries, with particular emphasis on reconfigurable intelligent surfaces (RIS). We model the radiation process as a compact operator acting between admissible source and observation spaces and characterize its symmetry through group equivariance. This formulation enables a systematic decomposition of the operator into irreducible representation sectors associated with the effective symmetry group, defined as the intersection of symmetries supported jointly by the source architecture, RIS geometry and programmability, receiver configuration, and propagation environment. We show that the FMM phase transition persists within each symmetry sector and that the total DoF budget is redistributed across sectors according to symmetry constraints. A key outcome of this analysis is the distinction between physical and effective degrees of freedom. While breaking the maximal $SO(2)$ symmetry does not increase the total number of electromagnetic DoF dictated by physics, symmetry reduction modifies their allocation across sectors, potentially lifting degeneracies and increasing the number of degrees of freedom that can be effectively addressed by a given excitation, RIS control, and measurement architecture, even when the total number of physical DoF remains fixed by fundamental limits. This clarifies the role of controlled symmetry breaking as a design mechanism rather than a means to surpass fundamental limits. The proposed framework bridges electromagnetic operator theory, representation theory, and RIS-enabled system design, providing both rigorous symmetry-resolved DoF accounting and actionable insights for excitation, surface programmability, and measurement strategies under practical architectural constraints.

Keywords: symmetry-resolved analysis; phase transitions; radiation operators; electromagnetic degrees of freedom; representation theory of compact groups; reconfigurable intelligent surfaces (RIS); mode resolvability

0. Introduction

Recent sub-THz receiver architectures pursue modular integration and aim to exploit reconfigurable intelligent surfaces (RIS) to shape the propagation environment at a structural level. In this context, RIS-assisted control is increasingly used not only to adapt baseband signaling, but to actively modify the electromagnetic radiation operator itself, enabling agile beamforming, multi-beam generation, and contoured footprint patterns.

A persistent challenge in this paradigm is the mismatch between the high-dimensional control spaces enabled by large arrays and programmable surfaces, and the fundamentally limited number of electromagnetic degrees of freedom (DoF) supported by wave physics. RIS optimization is often formulated in terms of large numbers of tunable parameters, while the physically available DoF are constrained by wavelength, aperture size, and geometry. Understanding how architectural and symmetry constraints mediate this gap is therefore essential for principled RIS design.

The physical-limit perspective on electromagnetic DoF was put on a rigorous footing by Franceschetti–Migliore–Minero (FMM), who showed that radiation operators across a cut exhibit a sharp phase transition in their singular-value spectrum: singular values remain approximately constant up to a cutoff index proportional to the boundary size measured in wavelengths, and then decay rapidly beyond that point [1]. This phenomenon yields an intrinsic, geometry-driven bound on the number of independent spatial channels, independent of fading or statistical channel models. Related signal- space interpretations of DoF have also appeared in the MIMO literature [2] and in optical and electromagnetic operator analyses [3].

While the FMM theory is often presented as a purely physical result, it implicitly relies on symmetry. In the canonical two-dimensional circular geometry, the singular-value analysis is carried out in the basis of cylindrical harmonics, which are precisely the irreducible representations of the rotation group $SO(2)$. In this sense, the FMM phase transition can be interpreted as a symmetry-resolved statement in which all symmetry sectors are simultaneously accessible. However, this symmetry structure remains implicit and is not treated as an analytical or design variable.

Related symmetry-based modal decompositions have appeared in various areas of mathematical physics and wave theory, typically in the context of exact separation of variables or highly idealized geometries. However, these approaches do not address the redistribution of physical versus effective degrees of freedom under architectural and programmability constraints, nor do they provide a symmetry-resolved extension of the FMM phase-transition phenomenon considered here.

The central objective of this work is to make symmetry explicit and to treat it as a first-class structural property of electromagnetic radiation operators. We develop a symmetry-resolved operator-theoretic framework in which the propagation process is modeled as a compact equivariant operator between admissible source and observation spaces. The effective symmetry group is defined as the intersection of symmetries jointly supported by the source architecture, RIS geometry and programmability, receiver configuration, and propagation environment. Within this framework, the FMM phase transition persists within each symmetry sector, while the total DoF budget is redistributed across sectors according to symmetry constraints.

A key outcome of this analysis is a clear distinction between physical and effective degrees of freedom. Breaking symmetry does not increase the total number of electromagnetic DoF dictated by physics, but it redistributes them across symmetry sectors, potentially lifting degeneracies and increasing the number of DoF that can be effectively excited, separated, and estimated by a given system architecture. Symmetry reduction is therefore understood not as a means to surpass fundamental limits, but as a controlled design mechanism for trading physical redundancy against practical accessibility.

To the best of our knowledge, this is the first work that makes the symmetry structure underlying the Franceschetti–Migliore–Minero phase transition explicit and treats it as a controllable operator-theoretic property, rather than as an implicit consequence of idealized geometries.

The remainder of the paper is organized as follows. Section 1 introduces a unified foundational framework combining physical DoF, radiation operators, and symmetry. Section 2 discusses symmetry groups of practical relevance for RIS-enabled planar systems. Section 3 develops the symmetry-resolved DoF bounds and phase-transition results. Engineering implications and design methodologies are discussed in Section 4. Technical proofs and equivariant Weyl arguments underlying the corollaries are provided in the appendices.

1. Physical Degrees of Freedom, Radiation Operators, and Symmetry

This section establishes the foundational framework used throughout the paper. We first recall the physical degrees-of-freedom (DoF) concept and the phase transition identified by Franceschetti–Migliore–Minero (FMM). We then introduce a minimal operator model for electromagnetic radiation across a cut and formalize symmetry through group actions and equivariance. These elements are combined into a unified perspective that supports the symmetry-resolved analysis developed in subsequent sections.

1.1. Physical Degrees of Freedom and the FMM Phase Transition

Franceschetti–Migliore–Minero model electromagnetic coupling across a cut by a compact linear operator mapping equivalent source currents to observed fields. In canonical settings, the singular values of this operator exhibit a sharp transition: they remain approximately constant up to a cutoff index proportional to the boundary size measured in wavelengths, and decay rapidly beyond that point [1]. This transition yields an intrinsic bound on the number of independent spatial channels supported by the propagation physics.

In two-dimensional circular geometries, the analysis is carried out explicitly using cylindrical harmonics $e^{jm\theta}$, which diagonalize the radiation operator and reveal the phase-transition structure. Although symmetry is not emphasized in the original formulation, this choice of basis already reflects the underlying rotational invariance of the geometry. The resulting DoF bound is therefore not merely a numerical coincidence, but a manifestation of the symmetry-resolved structure of the operator.

The present work builds on this observation by making symmetry explicit and extending the FMM viewpoint to scenarios in which symmetry is reduced, discrete, or deliberately controlled, as is typical in RIS-enabled systems.

1.2. Radiation Operators Across a Cut

Let $D \subset \mathbb{R}^d$ denote a source region and let O be an observation surface surrounding it. Let H_{in} denote a Hilbert space of admissible equivalent currents supported on ∂D (or on a suitable equivalent surface), and let H_{out} denote the Hilbert space of admissible observed fields on O . We model electromagnetic radiation across the cut by a linear operator

$$T : H_{\text{in}} \rightarrow H_{\text{out}}, \quad f \mapsto (E_t)|_O,$$

where E_t is the tangential electric field produced by the source f via the appropriate Green operator.

Under standard regularity and separation assumptions, such boundary-to-boundary operators are compact (often Hilbert–Schmidt), a property that underlies the very notion of a finite effective number of degrees of freedom [4–6]. Compactness guarantees the existence of a singular-value decomposition and provides a natural spectral framework for DoF analysis.

1.3. Group Actions, Equivariance, and Admissible Spaces

Let G be a finite or compact group acting on physical space by transformations that preserve geometry, material properties, boundary conditions, and the constitutive relations relevant to the radiation problem. Such transformations induce unitary representations $U_{\text{in}}(g)$ on H_{in} and $U_{\text{out}}(g)$ on H_{out} . The radiation operator T is said to be G -equivariant if

$$TU_{\text{in}}(g) = U_{\text{out}}(g)T, \quad \forall g \in G.$$

In RIS-enabled systems, not all source excitations or field patterns are admissible. Programmable surface states, hardware layouts, and measurement constraints restrict the admissible subspaces

$$H_{\text{in}}^{\text{adm}} \subseteq H_{\text{in}}, \quad H_{\text{out}}^{\text{adm}} \subseteq H_{\text{out}}.$$

These restrictions often enforce representation-theoretic constraints: only certain irreducible representations of G may be present in the admissible spaces. In this way, symmetry becomes a structural property of the operator and of the admissible signal spaces, rather than of individual signal realizations.

This framework provides a mathematically clean mechanism for modeling symmetry constraints and symmetry breaking induced by RIS programming. It sets the stage for the symmetry-resolved operator decomposition and degrees-of-freedom analysis developed in the following sections.

2. Symmetries of Interest in Sub-THz RIS Systems

The symmetry assumptions considered in this work are not introduced for mathematical convenience alone, but are directly motivated by the electromagnetic structure of metasurfaces and reconfigurable intelligent surfaces (RIS) operating at mmWave and sub-THz frequencies. It is well established that the geometric symmetry of both the unit-cell arrangement and the global aperture of a metasurface strongly constrains its admissible scattering responses, polarization properties, and modal degeneracies. These effects have been extensively studied in the context of electromagnetic metamaterials and metasurfaces, where symmetry groups such as cyclic and dihedral groups determine the allowed resonant modes and their coupling to incident fields [7].

Programmable and coding metasurfaces provide a direct physical precursor to RIS technology, demonstrating how discrete phase states and spatial arrangements can preserve or selectively break underlying geometric symmetries in a controlled manner [8]. At mmWave and sub-THz frequencies, RIS implementations inherit these symmetry constraints from their metasurface-based realizations, while adding reconfigurability as a design degree of freedom. As a result, RIS-assisted propagation environments naturally exhibit symmetry-induced degeneracies that depend on both the geometric layout of the surface and the chosen configuration of its programmable elements [9].

From a system-level perspective, these considerations motivate focusing on symmetry groups that arise naturally from rotationally symmetric layouts, periodic tilings, and mirror-symmetric arrangements commonly used in practical RIS and integrated antenna array designs. In particular, cyclic and dihedral groups provide representative and physically meaningful models for capturing the dominant symmetry constraints encountered in sub-THz RIS-assisted systems. These groups serve as the primary examples throughout this work, as they allow for a clear connection between geometric symmetry, operator equivariance, and the resulting sector-wise decomposition of the radiation operator.

(A) Planar lattice symmetries (approximate).

Many RIS are fabricated as periodic arrays of unit cells, motivating approximate invariance under translations by a lattice and point symmetries (e.g., C_4 , D_4). In idealized settings this connects to Bloch/Floquet decompositions and point-group irreps.

Accordingly, we focus on a small set of symmetry classes that arise naturally in practical RIS and array implementations, and that will serve as representative examples throughout the paper:

(B) Finite rotational/dihedral symmetries: C_n , D_n .

Polygonal/circular apertures, radomes, and symmetric packaging lead to n -fold rotations, possibly with reflections (D_n).

(C) Mirror symmetries: parity constraints.

Planar structures and symmetric feeds often induce one or multiple mirror planes, yielding an even/odd (parity) decomposition.

(D) Symmetry breaking by design/programming.

Off-axis beam steering and multi-beam patterns deliberately break the above symmetries; we model this as controlled transitions $G \rightarrow H$.

2.1. From RIS Symmetries to Symmetries of the Radiation Operator

For clarity, the discussion in this subsection focuses on RIS-induced symmetry constraints; additional symmetry-breaking effects arising from source and receiver configurations are addressed explicitly in Section 4.

The symmetries of a Reconfigurable Intelligent Surface (RIS), such as cyclic or dihedral symmetries (C_n , D_n), do not automatically translate into symmetries of the radiation operator T . Instead, they induce symmetries of T only insofar as they extend consistently to the full electromagnetic

configuration, including the admissible source space H_{in} , the observation space H_{out} , the surrounding medium, and the measurement setup.

Let G denote a group of geometric transformations acting on physical space. Symmetry $g \in G$ of the RIS is understood as an isometry preserving the RIS geometry and material properties, possibly together with a restricted class of programmable surface states. Such a transformation induces unitary representations

$$U_{\text{in}}(g) : H_{\text{in}} \rightarrow H_{\text{in}}, \quad U_{\text{out}}(g) : H_{\text{out}} \rightarrow H_{\text{out}},$$

defined by the natural pullback (and, when appropriate, vector transformation) of source currents and electromagnetic fields under the action of g .

As introduced before, the radiation operator $T : H_{\text{in}} \rightarrow H_{\text{out}}$ is said to be G -equivariant if it satisfies

$$T U_{\text{in}}(g) = U_{\text{out}}(g) T, \quad \forall g \in G.$$

This condition expresses that applying a symmetry transformation to the sources prior to radiation is equivalent to radiating first and then applying the same transformation to the observed field. Only under this equivariance condition can a symmetry of the RIS be regarded as a symmetry of the operator T .

Importantly, the effective symmetry group G is not determined by the RIS geometry alone, but by the entire electromagnetic configuration. Boundary conditions, background medium, admissible RIS programming patterns, and the geometry and sampling of the observation surface may all preserve or break a given symmetry. Consequently, a RIS with nominal D_n symmetry may lead to an operator T that is only C_n -equivariant, or even devoid of nontrivial symmetries, if reflections or rotations are broken by measurement or control constraints.

Example: C_n -symmetric RIS

To illustrate the symmetry-induced structure of the radiation operator, we consider a reconfigurable intelligent surface (RIS) with an n -fold rotational symmetry. The RIS consists of n identical unit cells arranged uniformly on a circle, so that a rotation by $2\pi/n$ maps the surface onto itself. This configuration induces a cyclic symmetry group C_n acting on the admissible signal spaces.

Let $\hat{C}$ denote the unitary operator implementing the action of the generator of C_n on the admissible space, i.e., the operator corresponding to a one-step cyclic rotation (or, equivalently, a permutation of equivalent RIS elements). Since the radiation operator T is equivariant with respect to this action, it commutes with $\hat{C}$ and with all operators generated by it.

As a consequence of this equivariance, the admissible space H admits a canonical decomposition into invariant subspaces associated with the irreducible representations of C_n ,

$$H = \bigoplus_{\alpha} V_{\alpha},$$

where α indexes the irreducible representations of C_n . In the present setting, these irreducible representations correspond to discrete angular harmonics compatible with the n -fold rotational symmetry. Each subspace V_{α} is invariant under both the group action and the radiation operator.

Accordingly, the radiation operator decomposes as a direct sum of independent blocks,

$$T = \bigoplus_{\alpha} T_{\alpha}, \quad T_{\alpha} := T|_{V_{\alpha}},$$

so that each symmetry sector can be analyzed independently. The dimension $\dim H_{\alpha}$ quantifies the number of electromagnetic degrees of freedom supported within the symmetry sector α . These sector-wise degrees of freedom are constrained by the same physical mechanisms underlying the FMM theory and collectively determine the total operator rank.

This decomposition is a direct consequence of the standard representation-theoretic structure of cyclic groups and does not rely on any assumption about the symmetry of individual excitations or RIS configurations; symmetry is a property of the operator and of the admissible space rather than of specific realizations (see, e.g., [10,11]). The resulting block-diagonal structure provides the foundation for the symmetry-resolved analysis developed in the subsequent sections.

Physically, this structure reflects a restriction of the angular spectrum: only angular modes whose indices are congruent modulo n can couple through the operator in the canonical circular harmonic picture. The discrete rotational symmetry of the RIS therefore enforces a block-diagonal structure on T , with independent symmetry sectors. If the RIS is programmed so as to break the C_n symmetry, additional couplings between sectors become possible, increasing the effective rank and expressivity of the operator. We further develop this example at system level in Section 4.

Remark. In the present C_n example, all irreducible representations are one-dimensional. Consequently, the abstract Schur factorization

$$T \simeq \bigoplus_{\alpha} (I_{V_{\alpha}} \otimes T_{\alpha})$$

reduces to a direct sum of operators acting on the corresponding multiplicity spaces. For notational simplicity, we denote these reduced operators by the same symbol T_{α} throughout this subsection.

3. Symmetry-Resolved Structure and Degrees of Freedom of the Radiation Operator

The goal of this section is to extend the physical degrees-of-freedom (DoF) phase-transition phenomenon identified by Franceschetti–Migliore–Minero to symmetry-constrained electromagnetic operators. To do so, we isolate a minimal set of structural assumptions that capture the essential physics of wave propagation across a cut while remaining agnostic to fine geometric or material details. These assumptions are not introduced for mathematical convenience alone: each reflects a property that is either inherent to electromagnetic radiation operators or directly motivated by the canonical FMM setting and its known generalizations.

Assumption 1 (Compactness and regularity). *The geometry and material configuration are such that the radiation operator is well-defined and compact (e.g., Hilbert–Schmidt) under the standard assumptions of electromagnetic scattering theory.*

Assumption 2 (FMM-type spectral envelope). *For each symmetry sector introduced below, the singular values admit an envelope with a sharp transition: there exist constants and a cutoff index scaling with the relevant boundary measure (in wavelengths) such that singular values are $O(1)$ below cutoff and decay rapidly above cutoff. This is the sector-wise analogue of the FMM phenomenon [1].*

Clarification. Assumption 2 is not an independent hypothesis, but a symmetry-resolved restatement of the Weyl-type spectral asymptotics underlying the original Franceschetti–Migliore–Minero theory. In particular, under standard regularity assumptions for electromagnetic cut operators, the existence of a sharp spectral transition and its sector-wise redistribution follow from equivariant Weyl splitting arguments, as made explicit in Appendix B.

Taken together, Assumptions 1–2 define a broad and physically meaningful class of operators. Compactness (Assumption 1) is a standard property of radiation operators across bounded cuts under mild regularity conditions, and underlies the very notion of a finite effective DoF. Assumption 2 does not require exact singular-value formulas or sharp cutoffs; rather, it posits the existence of a sector-wise spectral envelope with a transition scale consistent with the FMM physical limit. Partial violations of these assumptions—such as approximate equivariance, smooth rather than sharp spectral transitions, or weak symmetry-breaking perturbations—do not invalidate the qualitative conclusions. In such cases, the block-diagonal structure and the symmetry-resolved DoF bounds persist in an approximate sense, with constants and transition indices modified but with the same scaling behaviour.

3.1. Equivariant Operator Decomposition by Irreducible Representations

Let $\hat{G}$ denote the set of (equivalence classes of) irreps of G . Under Assumption 1, $\mathcal{H}_{\text{in}}$ and $\mathcal{H}_{\text{out}}$ decompose into isotypic components. For compact groups this follows from Peter–Weyl theory [12]; for finite groups it follows from classical representation theory [10,11]. Hence

$$\mathcal{H}_{\text{in}} \cong \bigoplus_{\alpha \in \hat{G}} V_{\alpha} \otimes \mathbb{C}^{m_{\alpha}^{\text{in}}}, \quad \mathcal{H}_{\text{out}} \cong \bigoplus_{\alpha \in \hat{G}} V_{\alpha} \otimes \mathbb{C}^{m_{\alpha}^{\text{out}}}, \quad (1)$$

where V_{α} is the carrier space of irrep α with dimension d_{α} , and $m_{\alpha}^{\text{in/out}}$ are multiplicities. Moreover, equivariance implies (via Schur’s lemma [10,11]) the operator decomposes as

$$T \cong \bigoplus_{\alpha \in \hat{G}} (\mathbf{I}_{V_{\alpha}} \otimes T_{\alpha}), \quad (2)$$

with $T_{\alpha} : \mathbb{C}^{m_{\alpha}^{\text{in}}} \rightarrow \mathbb{C}^{m_{\alpha}^{\text{out}}}$ compact.

3.2. Symmetry-Resolved Phase Transition and Degrees-of-Freedom Bound

Let $\{\sigma^{(\alpha)}\}$ denote the singular values of the sector-restricted operator $T^{(\alpha)}$, ordered in non-increasing order.

Theorem 1. Let $T : \mathcal{H}_{\text{in}} \rightarrow \mathcal{H}_{\text{out}}$ be a compact G -equivariant EM cut operator satisfying Assumptions 1–2, and let T_{α} be the symmetry blocks in (2). For any SNR level $\rho > 0$ and any accuracy threshold $\varepsilon \in (0, 1)$, define the effective number of DoF in sector α by

$$N_{\alpha}(\rho, \varepsilon) = \#\{k : \sigma_k^{(\alpha)}(T_{\alpha})^2 \geq \varepsilon/\rho\}, \quad (3)$$

and the total effective DoF as

$$N_{\text{eff}}(\rho, \varepsilon) = \sum_{\alpha \in \hat{G}} d_{\alpha} N_{\alpha}(\rho, \varepsilon). \quad (4)$$

Then:

(i) **Sector-wise phase transition.** For each α , the singular values of T_{α} exhibit a sharp transition at an index

$$k_c^{(\alpha)} \asymp \frac{B_{\alpha}}{\lambda^{d-1}}, \quad (5)$$

where B_{α} is a symmetry-dependent effective boundary measure and λ is the wavelength, consistent with the FMM scaling of the full operator [1]. Consequently, $N_{\alpha}(\rho, \varepsilon) = \Theta(k_c^{(\alpha)})$ for SNR not exponentially small.

(ii) **Symmetry-constrained DoF upper bound.** If the admissible subspaces $\mathcal{H}_{\text{in}}^{\text{adm}}$ and/or $\mathcal{H}_{\text{out}}^{\text{adm}}$ contain only irreps in a subset $\hat{G}_{\text{adm}} \subseteq \hat{G}$, then

$$N_{\text{eff}}^{\text{adm}}(\rho, \varepsilon) \leq \sum_{\alpha \in \hat{G}_{\text{adm}}} d_{\alpha} N_{\alpha}(\rho, \varepsilon) \lesssim \sum_{\alpha \in \hat{G}_{\text{adm}}} d_{\alpha} \frac{B_{\alpha}}{\lambda^{d-1}}. \quad (6)$$

(iii) **Information-theoretic upper bound (symmetry-resolved).** Let $C(\rho)$ denote the mutual-information capacity across the cut under average power/SNR ρ and optimal signaling within the admissible subspaces. Then

$$C(\rho) \leq \sum_{\alpha \in \hat{G}_{\text{adm}}} d_{\alpha} \sum_{k \geq 1} \log(1 + \rho \sigma_k^2(T_{\alpha})) \leq N_{\text{eff}}^{\text{adm}}(\rho, \varepsilon) \log(1 + \rho/\varepsilon) + O(1), \quad (7)$$

which mirrors the signal-space DoF interpretation of capacity in wave-based MIMO models [2]. Here capacity is used in the standard signal-space sense of Poon–Foschini–Tse, as an upper bound induced by resolvable singular modes, not as an operational network capacity claim.

(iv) **Controlled symmetry breaking increases DoF monotonically.** Suppose RIS programming changes the symmetry from G to a subgroup H (residual symmetry), and correspondingly enlarges the admissible set of irreps. Then the effective DoF and the capacity upper bound are monotone non-decreasing, reflecting enlargement of the feasible signaling subspace [6].

Remark 1. Here the signal-to-noise ratio (SNR) is a system-level quantity, defined as the ratio between total transmitted power and receiver noise variance, and used to assess mode resolvability. The logarithmic dependence on SNR/ε in Theorem 1 follows the standard Gaussian signaling framework, but is employed here to quantify the number of resolvable singular modes of the radiation operator under a finite accuracy threshold, rather than to characterize asymptotic channel capacity. In particular, such operators exhibit exponentially or super-algebraically decaying singular values, so that the number of singular modes satisfying $\sigma^2 \gtrsim \varepsilon/\text{SNR}$ scales logarithmically with the inverse threshold. Each such resolvable singular mode corresponds to an independent degree of freedom that can be stably excited and estimated at the prescribed resolution, so that the factor $\log(1 + \text{SNR}/\varepsilon)$ directly quantifies the effective number of degrees of freedom accessible under a finite error tolerance. A detailed justification of this logarithmic scaling, based on the decay properties of the singular values of compact radiation operators and the resulting thresholded mode count, is provided in Appendix A.

3.3. Connection with the Canonical Two-Dimensional FMM Setting

The sector-wise spectral envelope assumed in Theorem 1 admits a direct and explicit realization in the canonical circular geometry studied by Franceschetti–Migliore–Minero [1], which serves as a concrete reference case.

In the two-dimensional circular setting, the radiation operator across a circular cut admits an explicit singular-value decomposition in terms of cylindrical harmonics $e^{jm\theta}$, $m \in \mathbb{Z}$. These functions form the irreducible representations of the rotation group $SO(2)$, each of dimension one. The singular values $\{\sigma_m\}$ derived in [1] are given in closed form (involving Bessel or Hankel functions) and exhibit a sharp transition:

$$\sigma_m \approx \text{const}, \quad |m| \lesssim \kappa R, \quad \sigma_m \ll 1, \quad |m| \gtrsim \kappa R,$$

where $\kappa = 2\pi/\lambda$ is the wavenumber and R is the radius of the circular observation curve.

This *step-like* behaviour corresponds to the phase transition of the underlying electromagnetic radiation operator that yields the $O(kR)$ physical degrees-of-freedom bound, and underpins the capacity scaling laws reported by Franceschetti–Migliore–Minero, in which additional logarithmic factors originate from network-level information-theoretic considerations.

In the language of the present work, each angular index m corresponds to a distinct symmetry sector (irrep) of $SO(2)$, and the global singular spectrum is obtained by aggregating the one-dimensional sector contributions. The cutoff index $|m| \approx kR$ therefore plays the role of the sector-wise transition index $k_c^{(\alpha)}$ in Theorem 1, with B_α proportional to the arc-length measure associated with that sector.

This explicit circular example shows that the “sector-wise spectral envelope” assumption is not an abstract hypothesis, but a direct generalization of the exact singular-value behaviour computed in [1]. The contribution of the present work is to make this structure explicit and transferable: for more general geometries and for discrete symmetry groups (e.g., C_n , D_n), the same physical mechanism persists, with symmetry enforcing a redistribution and sampling of angular-momentum modes across irreducible sectors.

It is useful at this point to clarify the interpretation and scope of the quantities d_α and B_α appearing in Theorem 1 and in the symmetry-resolved connection with FMM. In the present framework, d_α denotes the dimension of the carrier space of the irrep associated with the irreducible representation α of the effective symmetry group and accounts for the multiplicity with which the sector-wise singular values of T_α appear in the global spectrum. As such, d_α is determined by the structure of the symmetry group, which fixes the sectorization, together with the choice of admissible signal spaces on which the operator acts, which determines how these sectors are realized and populated in a given physical configuration, including effects of dimensionality, geometry, and boundary conditions.

By contrast, the constants B_α encode problem-specific calibration information: they depend on the particular realization of the radiation operator, including propagation conditions, boundary conditions, and the functional representation used to express the fields (e.g., Bessel–Hankel expansions in the canonical circular free-space case). While d_α controls how many degrees of freedom are available in each symmetry sector, B_α determines where the corresponding sector-wise phase transition occurs along the singular-value spectrum. This separation delineates the scope of the results established here: the symmetry-resolved structure and scaling laws are universal at the operator level, whereas their numerical calibration is representation- and geometry-dependent.

To summarize: In the circular SO(2)-symmetric case, each angular index already defines an independent symmetry sector; the present framework makes this implicit sectorization explicit and transferable to reduced and discrete symmetry groups.

Remark. In the symmetry-resolved reinterpretation of the Franceschetti–Migliore–Minero framework developed above, the constants B_α arise as sector-wise boundary measures obtained by equivariant redistribution of the classical FMM Weyl term. They are not specific to Theorem 1, but reflect the underlying Weyl-type asymptotics of the radiation operator under symmetry constraints, as detailed in Appendix B.

3.4. Corollaries for Relevant Symmetry Groups

The following corollaries do not introduce new physical limits on electromagnetic degrees of freedom. Rather, they make explicit how the existing Franceschetti–Migliore–Minero bound is redistributed across symmetry sectors, according to the group-dependent multiplicities d_α and the sector-wise calibration constants B_α arising from the symmetry-resolved formulation developed above. In particular, they illustrate how the general operator-theoretic structure specializes to symmetry groups commonly encountered in RIS-enabled planar systems.

Corollary 1 (C_n rotational symmetry: $1/n$ DoF sampling). *Consider a planar setting where the cut operator is approximately equivariant under C_n rotations, and admissible RIS patterns are restricted to a single symmetry sector, namely the C_n -invariant one, which we simply denote by $N^{\text{inv}}(\rho, \epsilon)$ for later reference.*

Then only one C_n irrep contributes, hence

$$N_{\text{eff}}^{\text{inv}}(\rho, \epsilon) \lesssim \frac{1}{n} \frac{B}{\lambda} \quad (d = 2), \quad (8)$$

up to geometry-dependent constants, i.e., the accessible DoF are reduced by a factor $\approx n$ relative to unconstrained excitation, consistent with the FMM physical DoF scaling when all sectors are accessible [1].

Corollary 2 (D_n dihedral symmetry: parity & reflection selection). *If the residual symmetry includes reflections (D_n), irreps split into reflection-even and reflection-odd sectors [10]. If the hardware/measurement chain suppresses one parity (common in symmetric feed networks), then the DoF bound improves by an additional factor ≈ 2 relative to the C_n -invariant case.*

Corollary 3 (Mirror symmetry: even/odd halving). *If the configuration is equivariant under a mirror group and only even (or only odd) patterns are admissible, then $N_{\text{eff}}^{\text{adm}}(\rho, \epsilon) \lesssim \frac{1}{2} B / \lambda^{d-1}$.*

4. Engineering Implications and Design Methodology

The results developed in Sections 2 and 3 characterize the radiation operator T under a symmetry group G through equivariance and symmetry-resolved spectral structure. In practical deployments, however, this group should be understood as the *effective* symmetry group of the full system, resulting from the joint action of the source configuration, RIS geometry and programming constraints, receiver architecture, and propagation environment. In the following, we therefore interpret the group G appearing in the theoretical analysis as the effective group G_{eff} induced by the physical deployment. This perspective allows us to translate the symmetry-resolved DoF bounds and phase-transition results

into concrete design guidelines: symmetry is treated neither as an idealized assumption nor as an inherent limitation, but as a tunable structural property that governs the trade-off between operator expressivity, achievable degrees of freedom, and design complexity. The subsections below articulate how system-level symmetries arise, which design parameters can enforce or relax them, and how controlled symmetry breaking can be used strategically when DoF limitations become performance-limiting. Section 4 provides finite-dimensional illustrations of the abstract operator-theoretic results established above.

4.1. System-Level Symmetries as Transformations of Admissible Signal Spaces

From an engineering viewpoint, symmetry is not a property of individual signals, but of the system architecture that generates and processes them. In particular, symmetry becomes relevant whenever different physical configurations or control actions lead to equivalent electromagnetic input-output behavior. Typical examples include permuting identical antenna elements, rotating a regular array, or reindexing equivalent programmable states of a RIS.

In this work, we formalize this idea by associating a symmetry group with each subsystem:

- G_s for the source architecture,
- G_{RIS} for the RIS architecture,
- G_r for the receiver or observation architecture.

Each group represents transformations that are physically meaningful for the corresponding subsystem, such as hardware reconfigurations, relabeling of equivalent elements, or admissible control operations.

Crucially, these transformations do not act on individual signal realizations, but on the *spaces of signals that the system can generate or process*. For the source, this space is denoted by H_{in} and consists of all excitation vectors that can be produced by the transmit architecture (e.g., port excitations of an antenna array). Similarly, H_{out} denotes the space of all signals that can be observed or reconstructed by the receiver.

Every physical transformation $g \in G_s$ induces, by construction, a linear operator

$$U_{\text{in}}(g) : H_{\text{in}} \rightarrow H_{\text{in}},$$

which maps any admissible excitation to the excitation obtained after applying the corresponding hardware transformation. For instance, permuting identical antenna elements induces a permutation matrix acting on excitation vectors, while rotating a regular array induces a unitary transformation corresponding to index shifting or phase rotation. An analogous construction applies at the receiver, where each $g \in G_r$ induces a linear operator

$$U_{\text{out}}(g) : H_{\text{out}} \rightarrow H_{\text{out}}.$$

In this sense, the term *admissible* refers simply to signals that the system architecture can physically generate or measure. No symmetry is assumed for any particular excitation or received field; symmetry is a property of how the architecture transforms signals as a whole. Mathematically, this corresponds to each group acting through a (typically unitary) representation on the corresponding signal space.

The radiation operator $T : H_{\text{in}} \rightarrow H_{\text{out}}$ is said to be equivariant under a transformation g if the induced actions on input and output spaces satisfy

$$T U_{\text{in}}(g) = U_{\text{out}}(g) T.$$

Only transformations that admit compatible actions at the source, RIS, and receiver lead to such equivariance. The effective symmetry group governing the operator is therefore given by the intersection

$$G_{\text{eff}} = G_s \cap G_{\text{RIS}} \cap G_r,$$

where the intersection is taken over transformations that can be applied consistently to all subsystems.

4.1.1. Worked Example: From Geometric Symmetries to Matrices Acting on Signals

This subsection provides a concrete, engineer-oriented example showing how a geometric symmetry group G induces explicit linear operators $U_{\text{in}}(g)$ and $U_{\text{out}}(g)$ acting on finite-dimensional signal vectors. The goal is to make the construction operational: given a symmetric hardware layout, we write down the corresponding matrices and verify equivariance at the operator level.

Geometry and Coordinates

Consider a 2D free-space setting with a circular transmitter (TX) array of N identical antenna elements placed on a circle of radius R_{tx} centered at the origin. Using polar coordinates, the element locations are

$$\mathbf{r}_k = (R_{\text{tx}}, \theta_k), \quad \theta_k = \frac{2\pi k}{N}, \quad k = 0, 1, \dots, N-1.$$

Assume the receiver (RX) is also a circular array with M identical elements on a circle of radius R_{rx} , with sampling angles

$$\boldsymbol{\rho}_m = (R_{\text{rx}}, \varphi_m), \quad \varphi_m = \frac{2\pi m}{M}, \quad m = 0, 1, \dots, M-1.$$

(For simplicity we omit polarization; the same construction applies componentwise when vector fields are used.)

Admissible signal spaces.

The transmit architecture can independently feed each TX element with a complex weight. Thus the admissible input space is

$$H_{\text{in}} = \mathbb{C}^N,$$

and an excitation is a vector $\mathbf{x} = (x_0, \dots, x_{N-1})^\top$, where x_k is the complex excitation applied to the antenna at angle θ_k . Similarly, the admissible output space is

$$H_{\text{out}} = \mathbb{C}^M,$$

and a measurement is $\mathbf{y} = (y_0, \dots, y_{M-1})^\top$, where y_m is the sampled complex field at angle φ_m .

The Symmetry Group as a Physical Transformation

Because both arrays are regular and centered, the natural symmetry group is the cyclic rotation group. Let g denote a rotation by an angle Δ , with $\Delta = 2\pi/N$ for the TX array. Physically, applying g means rotating the entire TX array by Δ around the origin. Since the antennas are identical and equally spaced, this rotation does not change the hardware; it only *relabels* which element is considered “index k ”.

From $g \in C_N$ to a Matrix $U_{\text{in}}(g)$

The induced action on excitation vectors is obtained by following the relabeling: after a rotation by Δ , the element previously at θ_{k-1} occupies the position θ_k . Therefore, the excitation assigned to index k after rotation equals the previous excitation of index $k-1$ (modulo N). Hence

$$(U_{\text{in}}(g)\mathbf{x})_k = x_{k-1 \pmod{N}}.$$

In matrix form, $U_{\text{in}}(g)$ is the $N \times N$ cyclic permutation matrix

$$U_{\text{in}}(g) = \begin{bmatrix} 0 & 0 & \cdots & 0 & 1 \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{bmatrix}, \quad U_{\text{in}}(g)^H U_{\text{in}}(g) = I_N,$$

so the action is unitary. Repeating the rotation N times returns to the original labeling, so $U_{\text{in}}(g)^N = I_N$, matching the group relation of C_N .

Analogous Construction at the Receiver

A rotation of the RX array by $2\pi/M$ induces the same type of action on samples:

$$(U_{\text{out}}(h)\mathbf{y})_m = y_{m-1 \pmod{M}},$$

where h is the RX rotation generator and $U_{\text{out}}(h)$ is the $M \times M$ cyclic permutation matrix.

A Simple Discrete Operator Model T

To keep the example concrete, approximate the radiation operator by a linear map

$$\mathbf{y} = T\mathbf{x}, \quad T \in \mathbb{C}^{M \times N},$$

where $T_{m,k}$ is the (complex) response at RX sample m due to TX element k (including path loss and phase). In free space with centered circular arrays, $T_{m,k}$ depends only on the angular difference $\varphi_m - \theta_k$ (and radii), so T has a *circulant-by-block* structure: shifting TX indices shifts RX responses accordingly.

Equivariance as a Checkable Matrix Identity

Equivariance of the operator under a common rotation means:

$$T U_{\text{in}}(g) = U_{\text{out}}(h) T,$$

i.e., permuting the TX excitation vector before propagation produces the same outcome as propagating first and then permuting the RX samples. For the cyclic permutation matrices above, this identity is equivalent to the intuitive statement that the system response is invariant under a consistent relabeling of identical, regularly spaced elements.

Intersection of Symmetries and the Effective Group

If $N \neq M$, not every TX rotation matches an RX rotation. The transformations that are consistent for both arrays form the effective symmetry group

$$G_{\text{eff}} = C_{\text{gcd}(N,M)}.$$

This is the practical meaning of “system-level symmetry”: only the rotations that can be applied consistently to *both* the input indexing and the output indexing yield an equivariance relation of the form $T U_{\text{in}} = U_{\text{out}} T$.

Connection to Sectors (Optional Intuition)

For cyclic groups, the eigenvectors of the permutation matrices are discrete Fourier modes. Thus, transforming to the DFT basis diagonalizes $U_{\text{in}}(g)$ and $U_{\text{out}}(h)$, and the operator T becomes approximately block-diagonal across Fourier indices. These blocks are the concrete manifestation of the symmetry *sectors* discussed in Sections V–VI.

Including RIS Symmetry and Explicit Group Intersection

We now extend the previous example by explicitly incorporating a RIS and its symmetry group, in order to make the intersection between source, RIS, and receiver symmetries fully explicit.

Consider a circular RIS of L identical unit cells placed on a circle of radius R_{RIS} , centered at the origin, with angular positions

$$\boldsymbol{\zeta}_\ell = (R_{\text{RIS}}, \psi_\ell), \quad \psi_\ell = \frac{2\pi\ell}{L}, \quad \ell = 0, 1, \dots, L-1.$$

Each RIS element applies a complex reflection coefficient α_ℓ to the incident field. The admissible RIS configurations can therefore be represented by vectors

$$\boldsymbol{\alpha} = (\alpha_0, \dots, \alpha_{L-1})^\top \in \mathbb{C}^L,$$

possibly subject to amplitude or phase constraints.

Because the RIS elements are identical and regularly spaced, the RIS architecture admits a cyclic symmetry group $G_{\text{RIS}} = C_L$. A rotation by $2\pi/L$ physically corresponds to relabeling the RIS elements. This induces a unitary operator

$$(U_{\text{RIS}}(r)\boldsymbol{\alpha})_\ell = \alpha_{\ell-1 \pmod{L}},$$

where r denotes the generator of C_L . As before, $U_{\text{RIS}}(r)$ is a cyclic permutation matrix satisfying $U_{\text{RIS}}(r)^L = I_L$.

System-Level Consistency of Transformations

We now have three symmetry groups:

$$G_s = C_N \quad (\text{source}), \quad G_{\text{RIS}} = C_L \quad (\text{RIS}), \quad G_r = C_M \quad (\text{receiver}).$$

A rotation by an angle Δ defines a valid symmetry of the radiation operator only if it can be applied consistently to all three subsystems, i.e., if it corresponds to an element of each group. This requires

$$\Delta = \frac{2\pi k}{N} = \frac{2\pi\ell}{L} = \frac{2\pi m}{M}$$

for some integers k, ℓ, m .

The set of such rotations forms the effective symmetry group

$$G_{\text{eff}} = C_{\text{gcd}(N, L, M)}.$$

In other words, only rotations by multiples of $2\pi / \text{gcd}(N, L, M)$ simultaneously relabel source excitations, RIS coefficients, and receiver samples in a compatible manner.

Equivariance with RIS Included

Including the RIS, the radiation operator can be written schematically as

$$\mathbf{y} = T(\boldsymbol{\alpha}) \mathbf{x},$$

where $T(\boldsymbol{\alpha})$ denotes the input–output operator for a given RIS configuration. Equivariance under a common rotation $g \in G_{\text{eff}}$ means that

$$T(U_{\text{RIS}}(g)\boldsymbol{\alpha}) U_{\text{in}}(g)\mathbf{x} = U_{\text{out}}(g) T(\boldsymbol{\alpha})\mathbf{x}.$$

Physically, this identity states that rotating the source excitation, the RIS configuration, and the receiver sampling in a consistent way leaves the overall input–output relation unchanged.

Engineering Interpretation

This construction makes the role of symmetry intersection explicit. Even if the RIS has a high nominal symmetry (large L), the effective symmetry governing the operator is reduced if the source or receiver supports fewer equivalent transformations. Conversely, increasing symmetry at the source or receiver can unlock additional symmetry sectors that the RIS is already capable of supporting.

This example shows concretely how symmetry groups associated with the source, RIS, and receiver combine through intersection, and how the resulting effective group determines the symmetry-resolved structure exploited in the remainder of the paper.

4.2. Physical Versus Effective Degrees of Freedom Under Symmetry Constraints

The symmetry framework introduced in the previous subsection characterizes the radiation operator at a structural level. In particular, for a given effective symmetry group G_{eff} , the operator admits a block-diagonal decomposition across irreducible representation (irrep) sectors, and the FMM theory provides a physical upper bound on the number of degrees of freedom (DoF) supported within each sector. These DoF are intrinsic to the electromagnetic operator and depend only on physical parameters such as wavelength, system size, and propagation environment.

However, the existence of physical DoF does not imply that all of them can be reliably exploited by a practical system. In realistic RIS-assisted architectures, excitation, control, and measurement constraints limit how many of the physically available modes can be effectively accessed. This motivates a distinction between *physical* and *effective* degrees of freedom.

Physical DoF refer to the dimension of the signal subspace supported by the radiation operator itself, as quantified by the FMM phase transition and its symmetry-resolved extensions. These DoF represent a fundamental limit: they cannot be increased by architectural choices or signal processing, and breaking symmetry does not create new physical modes.

Effective DoF, on the other hand, quantify how many independent signal dimensions can be practically excited, separated, and estimated by the combined source, RIS, and receiver architecture. Effective DoF depend on factors such as excitation diversity, RIS programmability, measurement geometry, conditioning of the input–output operator, and noise. As a result, effective DoF are generally lower than the corresponding physical DoF and may vary significantly with system design.

Symmetry plays a dual role in this context. High symmetry leads to a structured organization into symmetry sectors, each with a limited number of physical DoF. While this organization is optimal from a purely physical standpoint, it can be detrimental in practice if only a small subset of sectors can be adequately excited or observed. In such cases, many physical DoF remain inaccessible, and the effective DoF are substantially lower than the physical limit.

Reducing symmetry modifies this balance. When the effective symmetry group is reduced, previously independent sectors are merged, degeneracies are lifted, and coupling between modes is introduced. This process may reduce the total number of symmetry sectors and, in some cases, reduce the total physical DoF supported by the operator. Nevertheless, the resulting operator may exhibit improved conditioning and stronger coupling between accessible modes, enabling a larger fraction of the remaining physical DoF to be effectively exploited. Consequently, controlled symmetry reduction can increase effective DoF even when physical DoF decrease.

This distinction is central to the design implications discussed in the remainder of this section. Rather than attempting to maximize physical DoF—which are fixed by the propagation physics—the proposed framework enables symmetry to be used as a design parameter that trades physical redundancy for practical accessibility. The irrep-domain design principles developed in the next subsection build directly on this idea by identifying how symmetry sectors can be selectively enforced or merged to optimize effective system performance.

4.3. A Simple Engineering Example: Symmetry Breaking Increases Effective DoF via Conditioning

The distinction between physical and effective degrees of freedom can be made concrete with a minimal example that requires no representation theory. The key mechanism is numerical conditioning: under highly symmetric architectures, many physically supported modes become nearly colinear at the receiver, so they cannot be stably separated in noise. A small, controlled symmetry breaking can lift this degeneracy and increase the number of reliably usable (effective) DoF without changing the underlying physical limits. As a minimal finite-dimensional illustration of the radiation operator T , we consider the following stylized matrix model.

Toy Model: A 4-Element Transmitter and a 2-Probe Receiver

Consider a transmitter with four identical radiating elements placed symmetrically (e.g., along a line with symmetric positions about the origin, or on a small symmetric aperture). Let the receiver measure the field at two probes (or two effective observation modes), producing a linear measurement model

$$y = Tx + w,$$

where $x \in \mathbb{C}^4$ collects complex excitation weights of the four elements, $y \in \mathbb{C}^2$ are the measured samples, $T \in \mathbb{C}^{2 \times 4}$ is the transfer matrix induced by propagation and the measurement setup, and w is additive noise.

Physical degrees of freedom are constrained by the physics of radiation and by the aperture sizes (as in the FMM setting), but even within those limits, the *effective* number of degrees of freedom depends on the conditioning of the mapping $x \mapsto y$. In practice, a convenient proxy is the number of singular values of T above the noise/accuracy threshold.

Case A: High Symmetry—Robustness Through Mode Concentration

We first consider a highly symmetric configuration, where the transmitter geometry and excitation respect all symmetries of the underlying architecture. In this regime, multiple physically admissible excitation patterns generate nearly identical field distributions at the receiver, leading to strong degeneracies in the input–output mapping.

A minimal stylized instance capturing this behavior is the following measurement model:

$$T_{\text{sym}} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix},$$

where the columns represent the field contributions of four symmetric transmit elements, and the rows correspond to two measurement probes. Although four independent excitations are physically available at the transmitter, all of them induce essentially the same response at the receiver. Algebraically, T_{sym} has rank one and a single nonzero singular value.

From an engineering perspective, this rank deficiency is not merely a limitation but also a source of robustness. Because transmitted energy and information are concentrated into a single dominant singular mode, noise and modeling errors are not significantly amplified. Small calibration errors, geometric perturbations, or mismatches between nominal and actual hardware configurations produce only minor changes in the observable response, since all admissible excitations project onto nearly the same measurement direction. In practice, this corresponds to a highly stable but weakly expressive system.

The drawback of this configuration is a reduced number of effective degrees of freedom. While the underlying physics supports multiple excitation patterns, symmetry-induced degeneracy collapses them into a single usable channel. As a result, multiplexing capability and spatial selectivity are limited: only one independent data stream or control mode can be exploited reliably. This loss of expressiveness is the price paid for the enhanced robustness provided by the symmetric design.

Case B: Controlled Symmetry Breaking—Increased Expressiveness at the Cost of Robustness

We now relax the symmetry of the previous configuration in a controlled manner. Rather than redesigning the architecture entirely, a minimal symmetry-breaking perturbation is introduced so that previously degenerate excitation patterns become distinguishable at the receiver. This can be achieved, for instance, by applying different phase shifts to subsets of transmit elements, introducing small geometric offsets, or selecting distinct RIS states over symmetric regions.

A corresponding stylized measurement model is

$$T_{\text{br}} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \end{bmatrix},$$

which differs from T_{sym} only in the response of the second measurement probe. Physically, this perturbation allows the receiver to discriminate between two previously indistinguishable excitation patterns. Algebraically, the rank of the input–output mapping increases to two, and two singular values of comparable magnitude emerge.

From an engineering perspective, this symmetry breaking increases the number of effective degrees of freedom: two independent excitation modes can now be reliably exploited, enabling spatial selectivity or multiplexing that was impossible in the fully symmetric case. This increased expressiveness is the primary benefit of relaxing symmetry constraints, as it allows the system to access a richer set of input–output behaviors without changing the underlying physical aperture or operating frequency.

However, this gain in effective degrees of freedom comes at a cost. Because the additional singular mode is typically weaker than the dominant one, the system becomes more sensitive to noise, calibration errors, and modeling inaccuracies. Small perturbations that were previously absorbed by symmetry now directly affect performance along the newly activated mode. In practice, this implies tighter calibration requirements, reduced noise margins, and increased control complexity.

Although this example is deliberately simple, it illustrates the central engineering trade-off underlying symmetry-aware design. Enforcing symmetry concentrates energy into a small number of robust modes, while controlled symmetry breaking redistributes that energy across multiple distinguishable modes. The optimal operating point is therefore not at either extreme, but at a balance where enough symmetry is preserved to ensure stability and robustness, while enough symmetry is relaxed to activate the desired number of effective degrees of freedom.

Interpretation in RIS Terms

In RIS-assisted systems, the stylized behavior illustrated above has a direct physical interpretation. A highly symmetric RIS configuration, where all elements share the same or symmetry-related phase states, effectively maps a large set of admissible excitations into nearly colinear radiation patterns, mirroring the rank-deficient behavior of H_{sym} . This concentration yields strong robustness but limits expressiveness, as multiple physical modes collapse into a small number of symmetry-adapted responses. By contrast, programming the RIS with distinct phase states over selected subsets, or deliberately enforcing only a subgroup of the full geometric symmetry, introduces controlled asymmetries that separate previously degenerate responses. In operator terms, this lifts symmetry-induced degeneracies, activates additional singular modes above the operational threshold, and increases the number of effective degrees of freedom. Crucially, the RIS enables this trade-off to be tuned in a structured and reversible manner, allowing symmetry to be treated as a programmable design parameter rather than a fixed constraint.

4.4. Practical Identification and Use of Effective Degrees of Freedom

This subsection addresses how the symmetry-resolved operator structure translates into practically usable degrees of freedom under realistic excitation and measurement architectures. While the sector multiplicities d_α determine the maximum number of physical degrees of freedom supported

within each symmetry sector, practical usability is governed by how these degrees of freedom are excited, observed, and conditioned in a given system.

A central observation is that physical support alone does not guarantee practical accessibility. Degrees of freedom become operational only when the corresponding sector is simultaneously excitable, observable, and sufficiently well conditioned. This leads to the following guiding principle:

Effective degrees of freedom are determined not only by symmetry-supported operator structure, but by the interaction between that structure and the excitation and measurement architecture.

In particular, sector-dependent constants such as B_α , which encode problem-specific calibration effects related to geometry, propagation conditions, and boundary conditions, determine where the corresponding sector-wise phase transition occurs along the singular-value spectrum and thus whether a sector becomes effectively accessible. From an engineering perspective, symmetry sectors may therefore be classified as accessible, weakly accessible, or effectively inaccessible depending on excitation diversity, measurement resolution, and numerical conditioning.

Operationally, this corresponds to determining, within each sector, which singular modes of $T^{(\alpha)}$ lie above the noise or accuracy threshold defining $N_\alpha(\rho, \epsilon)$. Sector-wise singular-value decay and conditioning thus provide direct indicators of practical usability.

Sector-Wise Operator Decomposition

Given G_{eff} , the admissible input and output spaces decompose as

$$H_{\text{in}} = \bigoplus_{\alpha} H_{\text{in}}^{(\alpha)}, \quad H_{\text{out}} = \bigoplus_{\alpha} H_{\text{out}}^{(\alpha)},$$

where α indexes irreducible representations of G_{eff} . Correspondingly, the radiation operator admits a block-diagonal structure

$$T = \bigoplus_{\alpha} T^{(\alpha)}, \quad T^{(\alpha)} : H_{\text{in}}^{(\alpha)} \rightarrow H_{\text{out}}^{(\alpha)}.$$

Each block represents an independent symmetry sector and can therefore be analyzed separately when assessing effective degrees of freedom.

Excitability of a Sector

A symmetry sector is excitable if the source architecture can generate signals with non-negligible projection onto $H_{\text{in}}^{(\alpha)}$. In practice, this is assessed by projecting the set of admissible excitation vectors onto the sector. If this projection is trivial, the sector is not excitable and its physical degrees of freedom cannot be accessed. Sectors with very small projections are weakly excitable and typically contribute little to system performance.

Observability of a Sector

Even if a sector can be excited, it may not be observable at the receiver. Observability is determined by the action of $T^{(\alpha)}$ on the excited subspace. In practice, one evaluates the effective rank and singular-value distribution of the restricted operator. If the corresponding singular values fall below the noise floor or collapse onto a low-dimensional subspace, the sector is effectively unobservable.

Conditioning and Effective Usability

For sectors that are both excitable and observable, numerical conditioning becomes the decisive factor. Even when physical degrees of freedom exist, ill-conditioned blocks lead to unstable inversion, poor separation of signal components, and strong noise amplification. Poorly conditioned sectors are therefore effectively unusable despite being structurally supported by the operator.

Effective Degrees of Freedom

The effective degrees of freedom of the system can be defined as the sum, over all sectors, of the dimensions of well-conditioned observable subspaces within each $T^{(\alpha)}$. This quantity is generally

smaller than the total number of physical degrees of freedom predicted by FMM and depends explicitly on excitation diversity, RIS programmability, measurement geometry, and noise.

Implications for Symmetry Reduction

Sectors that are non-excitable, unobservable, or poorly conditioned represent physical degrees of freedom that the system cannot exploit. These sectors are natural candidates for controlled symmetry reduction. By reducing the effective symmetry group, multiple symmetry sectors may merge into fewer, higher-dimensional blocks, lifting degeneracies, enhancing modal coupling, and improving conditioning. In this sense, symmetry acts not only as a structural constraint but as a tunable parameter governing the transition from physical to effective degrees of freedom.

The previous considerations naturally induce a structured design procedure linking symmetry analysis to practical accessibility of degrees of freedom.

Practical Workflow

From an engineering standpoint, the symmetry-resolved framework suggests the following sequence:

- identify the effective symmetry group G_{eff} imposed by the system architecture;
- perform a sector-wise decomposition of the operator and quantify the corresponding physical degrees of freedom;
- assess sector accessibility by examining excitation capability, observability, and spectral conditioning;
- where accessibility is limited, selectively reduce symmetry to merge sectors and improve modal coupling;
- repeat the analysis until a satisfactory balance between physical support and practical usability is obtained.

This procedure emphasizes that effective degrees of freedom are not fixed solely by structural considerations, but emerge from the interplay between symmetry, excitation, and measurement.

5. Conclusions

This work has examined how symmetry structure shapes the transition from physical to effective degrees of freedom in radiation operators. While symmetry-supported multiplicities determine the maximum number of physically admissible modes, practical accessibility is governed by the interplay between operator structure, excitation capabilities, measurement geometry, and numerical conditioning.

The sector-wise formulation developed in this paper provides a natural framework for separating these roles. By decomposing the operator into irreducible symmetry sectors, one can distinguish between degrees of freedom that are structurally supported and those that remain practically usable. This perspective highlights that effective degrees of freedom should not be interpreted as purely intrinsic properties of the propagation operator, but rather as emergent quantities arising from the interaction between symmetry constraints and system architecture.

A key implication is that symmetry is not merely a descriptive attribute of the physical configuration. Instead, it acts as a controllable parameter that mediates the trade-off between theoretical limits and operational performance. Reducing symmetry, when appropriately guided by sector accessibility and conditioning considerations, can merge degenerate subspaces, enhance modal coupling, and ultimately increase the number of usable degrees of freedom despite a possible reduction in nominal physical multiplicities.

From a methodological standpoint, the analysis suggests a structured pathway for translating symmetry information into design decisions. Evaluating excitability, observability, and conditioning at the sector level enables a systematic identification of accessible subspaces and provides a principled basis for symmetry reduction strategies.

More broadly, the results emphasize that the gap between physical and effective degrees of freedom is not a limitation to be passively accepted, but a quantity that can be actively shaped through architectural choices. In this sense, symmetry-resolved operator analysis offers both a conceptual clarification and a practical tool for guiding the design of complex electromagnetic systems.

Future work may extend this perspective to adaptive architectures and dynamically reconfigurable platforms, where symmetry itself becomes part of the design space rather than a fixed constraint.

Funding: This research was partially funded by the Spanish Ministry of Science, Innovation and Universities under Project HARMONIC-AIRIS (Grant No. 4353146064-146064-4-823).

Appendix A. Proof of Theorem 1: Symmetry-Resolved Phase Transition and DoF Bound

In this appendix we present a self-contained derivation of Theorem 1, in the sense that all representation-theoretic steps required for the symmetry-resolved decomposition are stated explicitly, without assuming prior familiarity with Schur–Weyl duality.

Throughout, G is a finite or compact group, $\mathcal{H}_{\text{in}}, \mathcal{H}_{\text{out}}$ are separable Hilbert spaces, and $T : \mathcal{H}_{\text{in}} \rightarrow \mathcal{H}_{\text{out}}$ is a compact G -equivariant operator, i.e.,

$$T U_{\text{in}}(g) = U_{\text{out}}(g) T, \quad \forall g \in G, \quad (\text{A1})$$

for unitary representations $U_{\text{in}}, U_{\text{out}}$.

Appendix A.1. Isotypic Decomposition and Block Diagonalization

We first recall the standard isotypic decomposition of a unitary representation. For each $\alpha \in \widehat{G}$, let V_α denote a fixed carrier space of the irrep α , with dimension $d_\alpha := \dim(V_\alpha)$.

Lemma A1 (Isotypic decomposition). *Assume G is finite or compact and the representations $U_{\text{in}}, U_{\text{out}}$ are unitary. Then there exist Hilbert spaces $\mathcal{M}_\alpha^{\text{in}}, \mathcal{M}_\alpha^{\text{out}}$ (multiplicity spaces) such that*

$$\mathcal{H}_{\text{in}} \cong \bigoplus_{\alpha \in \widehat{G}} (V_\alpha \otimes \mathcal{M}_\alpha^{\text{in}}), \quad \mathcal{H}_{\text{out}} \cong \bigoplus_{\alpha \in \widehat{G}} (V_\alpha \otimes \mathcal{M}_\alpha^{\text{out}}), \quad (\text{A2})$$

and $U_{\text{in}}(g)$ and $U_{\text{out}}(g)$ act as

$$U_{\text{in}}(g) \cong \bigoplus_{\alpha \in \widehat{G}} (\rho_\alpha(g) \otimes \mathbf{I}_{\mathcal{M}_\alpha^{\text{in}}}), \quad U_{\text{out}}(g) \cong \bigoplus_{\alpha \in \widehat{G}} (\rho_\alpha(g) \otimes \mathbf{I}_{\mathcal{M}_\alpha^{\text{out}}}), \quad (\text{A3})$$

where ρ_α is the irrep α on V_α .

Proof. This is standard for finite groups (complete reducibility) and compact groups (Peter–Weyl theory). Unitary representations decompose as orthogonal direct sums of irreducibles with multiplicities, yielding (A2)–(A3). See [11]. $\square$

Equivariance (A1) forces T to preserve the isotypic components and to factor through the multiplicity spaces.

Lemma A2 (Schur block form). *Under the decomposition of Lemma A1, the equivariant operator T satisfies*

$$T \cong \bigoplus_{\alpha \in \widehat{G}} (\mathbf{I}_{V_\alpha} \otimes T_\alpha), \quad (\text{A4})$$

for bounded operators $T_\alpha : \mathcal{M}_\alpha^{\text{in}} \rightarrow \mathcal{M}_\alpha^{\text{out}}$. If T is compact, then each T_α is compact.

Proof. Let P_α^{in} and P_α^{out} be the orthogonal projections onto the α -isotypic components in $\mathcal{H}_{\text{in}}$ and $\mathcal{H}_{\text{out}}$, respectively. Equivariance implies $TP_\alpha^{\text{in}} = P_\alpha^{\text{out}}T$ and thus T maps each isotypic component into the corresponding one. Restrict T to a single component $V_\alpha \otimes \mathcal{M}_\alpha^{\text{in}}$ and use (A1)–(A3):

$$(\rho_\alpha(g) \otimes \mathbf{I}) T_\alpha^{\text{full}} = T_\alpha^{\text{full}} (\rho_\alpha(g) \otimes \mathbf{I}) \quad \forall g \in G.$$

By Schur's lemma, any intertwiner between ρ_α and itself is proportional to the identity on V_α , hence the restriction has the form $\mathbf{I}_{V_\alpha} \otimes T_\alpha$ for some T_α acting on multiplicities. Compactness passes to direct summands and tensor factors, hence each T_α is compact if T is. $\square$

Appendix A.2. Singular Values: Multiplicity and Union Across Sectors

Throughout this appendix, we write $\sigma_k^{(\alpha)} := \sigma_k(T_\alpha)$ for the k -th singular value of the sector-restricted operator T_α .

Let $\{\sigma_k(A)\}_{k \geq 1}$ denote the singular values of a compact operator A in non-increasing order, counted with multiplicity.

Lemma A3 (Singular values of $\mathbf{I} \otimes T_\alpha$). *Let $T_\alpha : \mathcal{M}_\alpha^{\text{in}} \rightarrow \mathcal{M}_\alpha^{\text{out}}$ be compact and $d_\alpha = \dim(V_\alpha) < \infty$. Then the nonzero singular values of $\mathbf{I}_{V_\alpha} \otimes T_\alpha$ are exactly the singular values of T_α , each repeated d_α times:*

$$\sigma(\mathbf{I}_{V_\alpha} \otimes T_\alpha) = \{\sigma_k(T_\alpha) \text{ with multiplicity } d_\alpha\}. \quad (\text{A5})$$

Proof. Compute

$$(\mathbf{I} \otimes T_\alpha)^*(\mathbf{I} \otimes T_\alpha) = \mathbf{I} \otimes (T_\alpha^* T_\alpha).$$

Eigenvalues of $\mathbf{I} \otimes (T_\alpha^* T_\alpha)$ are eigenvalues of $T_\alpha^* T_\alpha$ each repeated d_α times, and singular values are the square roots of these eigenvalues. $\square$

Lemma A4 (Singular values of a direct sum). *Let $A = \bigoplus_\alpha A_\alpha$ with each A_α compact. Then the multiset of singular values of A is the multiset union of singular values of the A_α , re-ordered non-increasingly.*

Proof. $A^* A = \bigoplus_\alpha A_\alpha^* A_\alpha$. The spectrum of $A^* A$ is the union of spectra of $A_\alpha^* A_\alpha$ (with multiplicities), hence the claim for singular values. $\square$

Combining Lemmas A2–A4, the global singular spectrum of T is obtained by collecting sector singular values $\sigma_k(T_\alpha)$ and repeating them d_α times.

Appendix A.3. Effective DoF Bound (Part (ii) of Theorem 1)

Fix SNR $\rho > 0$ and accuracy threshold $\varepsilon \in (0, 1)$. Define the sector effective DoF as in the theorem:

$$N_\alpha(\rho, \varepsilon) = \#\{k : \sigma_k(T_\alpha)^2 \geq \varepsilon/\rho\}, \quad N_{\text{eff}}(\rho, \varepsilon) = \sum_{\alpha \in \widehat{G}} d_\alpha N_\alpha(\rho, \varepsilon). \quad (\text{A6})$$

Lemma A5 (Truncation dimension bound). *Let A be compact with singular values $\sigma_k(A)$. For any threshold $\tau > 0$, the number of singular values satisfying $\sigma_k(A) \geq \tau$ equals the minimal rank r such that there exists an operator A_r of rank r with $\|A - A_r\| \leq \tau$ (operator norm), and A_r can be chosen as the truncated SVD of A .*

Proof. Standard Eckart–Young–Mirsky theorem for compact operators: the best rank- r approximation error in operator norm is $\sigma_{r+1}(A)$. Therefore the smallest r such that $\sigma_{r+1}(A) < \tau$ is precisely the count of singular values $\geq \tau$. $\square$

Now consider symmetry constraints. Suppose admissible subspaces select only a subset of irreps $\hat{G}_{\text{adm}} \subseteq \hat{G}$, i.e., signals are restricted to

$$\mathcal{H}_{\text{in}}^{\text{adm}} = \bigoplus_{\alpha \in \hat{G}_{\text{adm}}} (V_{\alpha} \otimes \mathcal{M}_{\alpha}^{\text{in}}), \quad \mathcal{H}_{\text{out}}^{\text{adm}} = \bigoplus_{\alpha \in \hat{G}_{\text{adm}}} (V_{\alpha} \otimes \mathcal{M}_{\alpha}^{\text{out}}).$$

Let T^{adm} denote the restriction of T to these subspaces (equivalently, compress T by the orthogonal projections). Then by Lemma A2,

$$T^{\text{adm}} \cong \bigoplus_{\alpha \in \hat{G}_{\text{adm}}} (\mathbf{I}_{V_{\alpha}} \otimes T_{\alpha}).$$

By Lemmas A3 and A4, the number of singular values of T^{adm} above threshold $\sqrt{\varepsilon/\rho}$ is exactly $\sum_{\alpha \in \hat{G}_{\text{adm}}} d_{\alpha} N_{\alpha}(\rho, \varepsilon)$, i.e.,

$$N_{\text{eff}}^{\text{adm}}(\rho, \varepsilon) \leq N_{\text{eff}}(\rho, \varepsilon) = \sum_{\alpha \in \hat{G}_{\text{adm}}} d_{\alpha} N_{\alpha}(\rho, \varepsilon), \quad (\text{A7})$$

where equality holds by definition of restriction.

Finally, Assumption (FMM-type spectral envelope) in the main text asserts that each sector exhibits a sharp transition at index $k_c^{(\alpha)} \asymp B_{\alpha} / \lambda^{d-1}$, so for SNR not exponentially small,

$$N_{\alpha}(\rho, \varepsilon) = \Theta(k_c^{(\alpha)}) \asymp \frac{B_{\alpha}}{\lambda^{d-1}}. \quad (\text{A8})$$

Substituting (A8) into (A7) yields the DoF upper bound claimed in Theorem 1(ii).

Appendix A.4. Capacity Upper Bound (Part (iii) of Theorem 1)

Consider the (cut) channel in diagonalized form using the SVD of T^{adm} : its singular values are those of $\mathbf{I} \otimes T_{\alpha}$ across $\alpha \in \hat{G}_{\text{adm}}$, i.e., $\sigma_k(T_{\alpha})$ repeated d_{α} times. Under standard Gaussian signaling and total SNR ρ , the mutual-information contribution of each singular mode is $\log(1 + \rho\sigma^2)$, hence

$$C(\rho) \leq \sum_{\alpha \in \hat{G}_{\text{adm}}} d_{\alpha} \sum_{k \geq 1} \log(1 + \rho \sigma_k(T_{\alpha})^2). \quad (\text{A9})$$

This establishes the first inequality of Theorem 1(iii).

To obtain the simplified bound in terms of $N_{\text{eff}}^{\text{adm}}(\rho, \varepsilon)$, split the sum into “large” and “small” modes:

$$\mathcal{I}_{\alpha}^{\text{large}} := \{k : \sigma_k(T_{\alpha})^2 \geq \varepsilon/\rho\}, \quad \mathcal{I}_{\alpha}^{\text{small}} := \{k : \sigma_k(T_{\alpha})^2 < \varepsilon/\rho\}.$$

For $k \in \mathcal{I}_{\alpha}^{\text{large}}$, we have $\log(1 + \rho\sigma_k^2) \leq \log(1 + \rho/\varepsilon)$. Therefore,

$$\sum_{\alpha} d_{\alpha} \sum_{k \in \mathcal{I}_{\alpha}^{\text{large}}} \log(1 + \rho\sigma_k(T_{\alpha})^2) \leq \left(\sum_{\alpha} d_{\alpha} |\mathcal{I}_{\alpha}^{\text{large}}| \right) \log(1 + \rho/\varepsilon) = N_{\text{eff}}^{\text{adm}}(\rho, \varepsilon) \log(1 + \rho/\varepsilon). \quad (\text{A10})$$

For the “small” tail, use $\log(1 + x) \leq x$ for $x \geq 0$:

$$\sum_{\alpha} d_{\alpha} \sum_{k \in \mathcal{I}_{\alpha}^{\text{small}}} \log(1 + \rho\sigma_k^2) \leq \sum_{\alpha} d_{\alpha} \sum_{k \in \mathcal{I}_{\alpha}^{\text{small}}} \rho\sigma_k(T_{\alpha})^2. \quad (\text{A11})$$

Since each T_{α} is compact and typically Hilbert–Schmidt under the EM operator assumptions, $\sum_k \sigma_k(T_{\alpha})^2 < \infty$. Moreover, on $\mathcal{I}_{\alpha}^{\text{small}}$ we have $\rho\sigma_k^2 < \varepsilon$, so the tail is bounded by a finite constant depending on $\{T_{\alpha}\}$ and ε (and does not scale with the DoF cutoff). Denoting this finite remainder by $O(1)$, (A9)–(A11) yield

$$C(\rho) \leq N_{\text{eff}}^{\text{adm}}(\rho, \varepsilon) \log(1 + \rho/\varepsilon) + O(1),$$

which is the second inequality in Theorem 1(iii).

Appendix A.5. Monotonicity Under Controlled Symmetry Breaking (Part (iv) of Theorem 1)

Let G be the initial symmetry and $H \subseteq G$ a residual symmetry after “breaking”. At the level of admissible signaling subspaces, symmetry breaking corresponds to enlarging the feasible subspaces:

$$\mathcal{H}_{\text{in}}^{(G)} \subseteq \mathcal{H}_{\text{in}}^{(H)}, \quad \mathcal{H}_{\text{out}}^{(G)} \subseteq \mathcal{H}_{\text{out}}^{(H)}.$$

Equivalently, the admissible irrep set expands from $\hat{G}_{\text{adm}}$ to $\hat{H}_{\text{adm}}$, activating additional sectors.

Since capacity is the supremum of mutual information over admissible input distributions, enlarging the admissible set cannot decrease it:

$$C^{(H)}(\rho) = \sup_{\text{inputs on } \mathcal{H}_{\text{in}}^{(H)}} I(\text{input}; \text{output}) \geq \sup_{\text{inputs on } \mathcal{H}_{\text{in}}^{(G)}} I(\text{input}; \text{output}) = C^{(G)}(\rho).$$

Likewise, N_{eff} as defined by counting modes above threshold is monotone under inclusion of admissible sectors: adding blocks contributes nonnegative counts $d_{\alpha} N_{\alpha}(\rho, \epsilon)$.

This proves Theorem 1(iv).

Remark.

Lemmas A1–A4 yield the symmetry block decomposition and the sector multiplicities in the singular spectrum, making explicit the underlying Schur–Weyl type duality between the representation-theoretic symmetry structure and the dynamical modal degrees of freedom. The DoF bound follows by threshold counting of singular values in each block, combined with the sector-wise phase transition envelope. The capacity bound is obtained by summing $\log(1 + \rho\sigma^2)$ over symmetry-resolved singular modes and bounding the contribution of modes above and below the resolution threshold. Monotonicity under symmetry breaking follows from inclusion relations between admissible signaling subspaces.

Appendix B. Equivariant Weyl Laws and Proofs of Corollaries 1–3

This appendix develops a unified spectral argument underlying the calibration constants appearing in Corollaries 1–3. We first establish a general equivariant Weyl splitting principle, in the spirit of classical equivariant Weyl laws, and then specialize it to the identified symmetry groups relevant to RIS-enabled planar systems [13,14].

Throughout, let T be the compact radiation (cut) operator acting on a separable Hilbert space H , and assume that T is equivariant under a finite group G with unitary representation $U(g)$, i.e.,

$$U(g)T = TU(g), \quad \forall g \in G.$$

Let $E(\mu)$ denote the spectral projector of the positive compact operator T^*T associated with the interval $[0, \mu]$, i.e., the orthogonal projection onto the subspace spanned by the singular vectors of T whose squared singular values are less than or equal to μ . Equivalently, if $\{\sigma_k(T)\}_{k \geq 1}$ are the singular values of T ordered non-increasingly, then $E(\mu)$ projects onto the span of those modes satisfying $\sigma_k(T)^2 \leq \mu$. The associated counting function

$$N(\mu) = \text{Tr } E(\mu)$$

therefore counts the number of singular modes of T below the threshold μ , or, equivalently, the number of eigenvalues of T^*T not exceeding μ .

We assume that, as the spectral threshold increases, $N(\mu)$ satisfies a Weyl asymptotic of the form, under the standard assumptions recalled above:

$$N(\mu) = c B \Phi(\mu) + o(\Phi(\mu)), \quad \Phi(\mu) \rightarrow \infty,$$

where B is the effective boundary measure and $c > 0$ is geometry-independent; and $\Phi(\mu)$ denotes a monotone scaling function capturing the universal growth rate of the singular-value counting function, independent of symmetry considerations. Its precise form depends on the analytic nature of the underlying operator and the spatial dimension, but is not required for the equivariant splitting arguments that follow; only the existence of a leading-order Weyl term with lower-order remainder is used.

We further assume the standard twisted-trace estimate: for every non-identity element $g \neq e$,

$$\text{Tr}(U(g)E(\mu)) = o(\Phi(\mu)).$$

This condition reflects the fact that nontrivial isometries have fixed sets of lower dimension and therefore cannot contribute to the leading Weyl term. Such estimates are classical for symmetric elliptic and pseudodifferential operators.

Appendix B.1. Equivariant Weyl Splitting

Let $\widehat{G}$ denote the set of irreducible representations of G . For each $\alpha \in \widehat{G}$ with character χ_α and dimension d_α , define the orthogonal isotypic projector

$$P_\alpha = \frac{d_\alpha}{|G|} \sum_{g \in G} \chi_\alpha(g^{-1}) U(g).$$

The sector-wise counting function is

$$N_\alpha(\mu) = \text{Tr}(P_\alpha E(\mu)), \quad N(\mu) = \sum_{\alpha \in \widehat{G}} N_\alpha(\mu).$$

Lemma A6 (Equivariant Weyl splitting). *Under the above assumptions,*

$$N_\alpha(\mu) = \frac{d_\alpha^2}{|G|} N(\mu) + o(\Phi(\mu)).$$

Proof. Using the projector formula and the commutation of T with the group action,

$$N_\alpha(\mu) = \frac{d_\alpha}{|G|} \sum_{g \in G} \chi_\alpha(g^{-1}) \text{Tr}(U(g)E(\mu)).$$

Separating the identity term gives

$$N_\alpha(\mu) = \frac{d_\alpha}{|G|} \chi_\alpha(e) \text{Tr}(E(\mu)) + \frac{d_\alpha}{|G|} \sum_{g \neq e} \chi_\alpha(g^{-1}) \text{Tr}(U(g)E(\mu)).$$

Since $\chi_\alpha(e) = d_\alpha$ and the twisted traces are $o(\Phi(\mu))$,

$$N_\alpha(\mu) = \frac{d_\alpha^2}{|G|} N(\mu) + o(\Phi(\mu)).$$

□

Interpreting the sector calibration constants B_α as the geometric coefficients of the leading Weyl term yields

$$B_\alpha = \frac{d_\alpha^2}{|G|} B.$$

Thus the leading-order boundary measure is redistributed uniformly across symmetry sectors according to representation dimension.

Appendix B.2. Cyclic Symmetry: Proof of Corollary 1

Let $G = C_n$. All irreducible representations are one-dimensional, so $d_\alpha = 1$ and $|G| = n$. The lemma immediately gives

$$N_\alpha(\mu) = \frac{1}{n}N(\mu) + o(\Phi(\mu)), \quad B_\alpha = \frac{B}{n}.$$

If admissible excitations are restricted to the invariant sector (the trivial character), the physical scaling $N_{\text{phys}} \asymp B/\lambda$ yields

$$N_{\text{eff}}^{\text{inv}} \asymp \frac{B}{n\lambda} = \frac{1}{n}N_{\text{phys}},$$

establishing the $1/n$ reduction in accessible degrees of freedom.

Appendix B.3. Dihedral Symmetry: Proof of Corollary 2

Let $G = D_n$. The rotational subgroup produces the same $1/n$ reduction as in the cyclic case. When reflections are included, the invariant rotational sector splits into even and odd parity subspaces.

Let s denote a reflection with unitary action $U(s)$. Define the parity projectors

$$P_\pm = \frac{1}{2}(I \pm U(s)).$$

The parity-resolved counting functions are

$$N_\pm(\mu) = \text{Tr}(P_\pm E(\mu)) = \frac{1}{2}N_{\text{inv}}(\mu) \pm \frac{1}{2}\text{Tr}(U(s)E(\mu)).$$

By the twisted-trace hypothesis, the second term is $o(\Phi(\mu))$, giving

$$N_\pm(\mu) = \frac{1}{2}N_{\text{inv}}(\mu) + o(\Phi(\mu)).$$

Consequently,

$$B_\pm = \frac{1}{2}B_{\text{inv}} = \frac{B}{2n}.$$

If hardware or measurement suppresses one parity, the effective operator is compressed to a single subsector, yielding

$$N_{\text{eff}}^{\text{parity}} \asymp \frac{B}{2n\lambda} = \frac{1}{2}N_{\text{eff}}^{\text{inv}},$$

which proves the additional factor ≈ 2 reduction.

Appendix B.4. Mirror Symmetry: Proof of Corollary 3

Consider the order-two group $G = C_s = \{e, s\}$ generated by a reflection. Since $U(s)^2 = I$, the spectrum lies in $\{+1, -1\}$ and the operator preserves parity, so

$$T \simeq T^{(+)} \oplus T^{(-)}.$$

Applying the equivariant splitting lemma with $|G| = 2$ gives

$$N_\pm(\mu) = \frac{1}{2}N(\mu) + o(\Phi(\mu)), \quad B^{(+)} = B^{(-)} = \frac{B}{2}.$$

Restricting admissible excitations to a single parity sector therefore yields

$$N_{\text{eff}}^{\text{adm}}(\rho, \epsilon) \asymp \frac{B}{2\lambda^{d-1}},$$

establishing the approximate halving of accessible degrees of freedom.

In summary, cyclic, dihedral, and mirror symmetry reductions all follow from the same equivariant redistribution of the leading Weyl term. The symmetry group determines how the boundary measure—and hence the physical degrees of freedom—is partitioned across irreducible sectors, while architectural constraints select which sectors remain effectively accessible.

References

1. Franceschetti, M.; Migliore, M.D.; Minero, P. The Capacity of Wireless Networks: Information-Theoretic and Physical Limits. *IEEE Transactions on Information Theory* **2009**, *55*, 3413–3424. <https://doi.org/10.1109/TIT.2009.2023705>.
2. Poon, A.S.Y.; Foschini, G.J.; Tse, D.N.C. Degrees of Freedom in Multiple-Antenna Channels: A Signal Space Approach. *IEEE Transactions on Information Theory* **2005**, *51*, 523–536. <https://doi.org/10.1109/TIT.2004.840894>.
3. Miller, D.A.B. Communicating with Waves between Volumes: Evaluating Orthogonal Spatial Channels and Limits on Coupling Strengths. *Applied Optics* **2000**, *39*, 1681–1699. <https://doi.org/10.1364/AO.39.001681>.
4. Jackson, J.D. *Classical Electrodynamics*, 3 ed.; John Wiley & Sons, 1999.
5. Chew, W.C. *Waves and Fields in Inhomogeneous Media*, 2 ed.; IEEE Press, 1995.
6. Conway, J.B. *A Course in Functional Analysis*, 2 ed.; Springer, 1990.
7. Lim, T.C. Metamaterials and Symmetry. *Symmetry* **2022**, *14*, 1587. <https://doi.org/10.3390/sym14081587>.
8. Cui, T.J.; Qi, M.; Wan, X.; Zhao, J.; Cheng, Q. Coding Metamaterials, Digital Metamaterials and Programmable Metamaterials. *Science* **2014**, *345*, 159–162. <https://doi.org/10.1126/science.1257571>.
9. Wu, Q.; Zhang, R. Towards Smart and Reconfigurable Environment: Intelligent Reflecting Surface Aided Wireless Network. *IEEE Communications Magazine* **2019**, *58*, 106–112. <https://doi.org/10.1109/MCOM.001.1900107>.
10. Serre, J.P. *Linear Representations of Finite Groups*; Springer, 1977.
11. Fulton, W.; Harris, J. *Representation Theory: A First Course*; Springer, 1991.
12. Folland, G.B. *A Course in Abstract Harmonic Analysis*; CRC Press, 1995.
13. Helffer, B. *Spectral Asymptotics in the Semi-Classical Limit*; Vol. 159, *London Mathematical Society Lecture Note Series*, Cambridge University Press, 1988.
14. Chavel, I. *Eigenvalues in Riemannian Geometry*; Academic Press: Orlando, FL, 1984.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.