
Journal of Computational Physics

Operator-adapted wavelets for finite-element differential forms

Max Budninskiy, Houman Owjadi, Mathieu Desbrun*

Caltech, 1200 E. California Blvd, MS 305-16, Pasadena, CA 91125, USA

ABSTRACT

We introduce in this paper an operator-adapted multiresolution analysis for finite-element differential forms. From a given continuous, linear, bijective, and self-adjoint positive-definite operator \mathcal{L} , a hierarchy of basis functions and associated wavelets for discrete differential forms is constructed in a fine-to-coarse fashion and in quasilinear time. The resulting wavelets are \mathcal{L} -orthogonal across all scales, and can be used to derive a Galerkin discretization of the operator such that its stiffness matrix becomes block-diagonal, with uniformly well-conditioned and sparse blocks. Because our approach applies to arbitrary differential p -forms, we can derive both scalar-valued and vector-valued wavelets block-diagonalizing a prescribed operator. We also discuss the generality of the construction by pointing out that it applies to various types of computational grids, offers arbitrary smoothness orders of basis functions and wavelets, and can accommodate linear differential constraints such as divergence-freeness. Finally, we demonstrate the benefits of the corresponding operator-adapted multiresolution decomposition for coarse-graining and model reduction of linear and non-linear partial differential equations.

Keywords: non- L^2 multiresolution analysis, finite-element differential forms, operator-adapted wavelets.

1. Introduction

The relevance of wavelets to the representation of integral and differential operators emerged very early on in the development of multiresolution analysis (MRA) [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11]. Wavelets have not only been used for the fast inversion of a given operator [12, 13] but also for its compression [14, 15, 16] and the analysis of solutions for the corresponding operator equations [17, 18].

Towards operator-adapted wavelets. While first generation adaptive wavelets (such as bi-orthogonal wavelets [19]) can be constructed with arbitrarily high regularity, their shift and scale invariance prevents from adapting them to irregular domains or non-homogeneous coefficients. This problem has stimulated the emergence of second generation

*Corresponding author: tel.: +1-626-395-6230; fax: +1-626-792-4257;

e-mail: mbudnins@caltech.edu (Max Budninskiy), owjadi@caltech.edu (Houman Owjadi), mathieu@caltech.edu (Mathieu Desbrun)

wavelets [20, 21, 22, 23, 24] offering stronger adaptability. While much more versatile, these wavelets have found greater adoption in signal processing than in other computational fields. In the context of numerical simulation, a consensus has now emerged that the ideal notion of “operator-adapted wavelets” should have three properties: they should be scale orthogonal with respect to the energy scalar product of the operator, the operator should be well-conditioned within each sub-band defined by those wavelets, and they should be localized in space. These properties are particularly desirable when wavelets are used as a basis of finite elements for the Galerkin discretization of the operator as the resulting stiffness matrix becomes block-diagonal, with uniformly well-conditioned and sparse blocks. Furthermore, these properties also imply that the wavelets can be used as Wannier functions as well [25, 26, 27] given their dual localization in space and eigenspace. Recently, scalar-valued operator-adapted wavelets satisfying these three properties for linear symmetric positive definite scalar differential operators have been formulated in [28, 29, 30, 31]. These so-called “gamblets” were derived from a game theoretic approach to numerical approximation [32], and were shown to be efficiently computable in a fine-to-coarse fashion in quasilinear complexity.

Contributions. In this paper, we adopt a finite-element perspective to the development of an operator-adapted multiresolution analysis. After formulating a series of axiomatic properties to ensure a proper Galerkin discretization tailored to a given continuous, linear, bijective, and self-adjoint positive-definite operator, we introduce a construction of operator-adapted basis functions and associated wavelets that matches the recent construction of gamblets introduced in [28]. However, our resolute finite element approach extends seamlessly from scalar-valued functions to *differential forms*, the building blocks of Cartan’s Exterior Calculus [33]. While Finite Element Exterior Calculus [34] and its underlying structure-preserving de Rham complex has been a recent topic of interest in computations, the construction of form-based wavelets in this context has received little to no attention, even if wavelets for discrete differential forms could be crucial to the discretization of linear and nonlinear partial differential equations (PDEs) arising in, for instance, electromagnetics and fluid dynamics. We thus propose a constructive approach to *multiresolution analysis of finite-element differential p -forms adapted to an exterior operator*. As a result, we introduce the first construction of vector-valued wavelets that block-diagonalize a given operator. We also show how to derive operator-adapted wavelets restricted to a linear subspace satisfying a differential constraint, such as divergence-free wavelets. We demonstrate that for time-dependent nonlinear systems (involving complex interactions between multiple scales and domain geometry), our wavelets can be computed on the fly in quasilinear complexity through linearization of the dynamics. In addition to our algorithmic presentation of the general construction of operator-adapted wavelets, we discuss how an operator-adapted multiresolution analysis provides many opportunities to construct reduced models or perform fast numerical integration.

2. From Classical to Operator-adapted Wavelet-based Finite Element Analysis

Before delving into our contributions, we provide background for our work by briefly discussing the usual wavelet-Galerkin approach, and recent proposals on how to improve upon its foundations.

2.1. Wavelet-Galerkin Approach

Let H be a subspace of the Sobolev space H^s (defined as the closure of smooth functions compactly supported in an open bounded domain Ω , with respect to the L^2 norm of total derivatives of order $s \in \mathbb{N}^*$), and let $\mathcal{L}: H \rightarrow H^* \subset H^{-s}$ be a continuous, linear, bijective, and self-adjoint positive-definite operator used in a differential equation of the form

$$\mathcal{L}u = g. \quad (1)$$

A standard procedure to analyze and solve Eq. (1) involves expressing it in a weak form through the introduction of a bilinear form (often called “energy scalar product”) $L: H^2 \rightarrow \mathbb{R}$ with $L(u, v) := \langle u, v \rangle_{\mathcal{L}} = \int_{\Omega} u \cdot \mathcal{L}v$ (where product operation represents pointwise inner product) and a linear operator $G: H \rightarrow \mathbb{R}$ as $G(v) := \langle g, v \rangle_{L^2} = \int_{\Omega} g \cdot v$. Since $\langle u, v \rangle_{\mathcal{L}}$ defines an inner product on H , the bilinear form is both continuous (i.e., $L(u, v) \leq C\|u\|_{\mathcal{L}}\|v\|_{\mathcal{L}}$) and coercive (i.e., $L(u, u) \geq C^{-1}\|u\|_{\mathcal{L}}^2$). As a result, the celebrated Lax-Milgram theorem implies that the weak problem

$$L(u, v) = G(v), \quad \forall v \in H \quad (2)$$

has a *unique* solution for any right-hand side $g \in H^*$ of Eq. (1).

The classical wavelet-Galerkin approach uses a compactly supported L^2 -orthogonal functional basis [35] to solve Eq. (2). It proceeds by approximating the subspace H by a finite-dimensional solution space \mathcal{V}^q and applying the multiresolution construction as follows. First, a nested sequence of lower-dimensional (coarser) functional spaces $\{\mathcal{V}^k\}_{k=1}^{q-1} \subset H$ are defined such that

$$\mathcal{V}^k \subset \mathcal{V}^{k+1} \text{ for } k = 1 \dots q-1.$$

Complementary wavelet spaces $\{\mathcal{W}^k\}_{k=1}^{q-1} \subset H$ are then derived by enforcing

$$\mathcal{V}^{k+1} = \mathcal{V}^k \oplus \mathcal{W}^k \text{ for } k = 1 \dots q-1, \quad (3)$$

where \oplus denotes the direct sum of L^2 -orthogonal subspaces. Thus, \mathcal{V}^q admits a multiresolution decomposition:

$$H \approx \mathcal{V}^q = \mathcal{V}^1 \oplus \mathcal{W}^1 \oplus \mathcal{W}^2 \oplus \dots \mathcal{W}^{q-2} \oplus \mathcal{W}^{q-1}$$

On each resolution level k (corresponding to \mathcal{V}^k), suppose we have n_k functions $\{\varphi_i^k\}_{i=1}^{n_k}$ forming a basis of \mathcal{V}^k and N_k (pre-)wavelets $\{\psi_j^k\}_{j=1}^{N_k}$ forming a basis of the wavelet space \mathcal{W}^k , with $N_k := n_{k+1} - n_k$ due to Eq. (3). (Throughout the paper, we will not require the wavelet basis elements of a fixed resolution k to be orthonormal; thus formally, $\{\psi_i^k\}_{i=1}^{N_k}$ are only pre-wavelets, although we will refer to them as wavelets for simplicity). The solution of Eq. (2) can now be approximated by a function u^q of the form:

$$u^q = \sum_{i=1}^{n_1} v_i^1 \varphi_i^1 + \sum_{k=1}^{q-1} \sum_{j=1}^{N_k} w_j^k \psi_j^k \quad (4)$$

Let $N := n_1 + \sum_{k=1}^{q-1} N_k = n_q$ be the total number of degrees of freedom associated with the wavelet decomposition of \mathcal{V}^q , and let \mathbf{w} be the vector of coefficients $(v^1, w^1, \dots, w^{q-1})$ of u^q of size N . The finite-dimensional weak problem can be formulated as a linear equation:

$$\mathbf{L}\mathbf{w} = \mathbf{g} \quad (5)$$

where \mathbf{L} is the global $N \times N$ stiffness matrix

$$\mathbf{L} = \begin{bmatrix} A^1 & M^{(1,2)} & \dots & M^{(1,q)} \\ M^{(2,1)} & B^1 & \dots & M^{(2,q)} \\ \vdots & \vdots & \ddots & \vdots \\ M^{(q-1,1)} & M^{(q,2)} & \dots & B^{q-1} \end{bmatrix}$$

composed of an $n_1 \times n_1$ submatrix $A_{ij}^1 := L(\varphi_i^1, \varphi_j^1)$ encoding the stiffness matrix of the coarsest basis functions, of a series of $N_k \times N_k$ submatrices $B_{ij}^k := L(\psi_i^k, \psi_j^k)$ that are the stiffness submatrices of wavelets for each resolution level, and the off-diagonal block matrices $M^{(r,s)}$ describing interactions between bases of different resolution levels r and s :

$$\begin{aligned} M_{ij}^{(1,s)} &:= L(\varphi_i^1, \psi_j^s) = M_{ji}^{(s,1)} \text{ for } s > 1, \\ M_{ij}^{(r,s)} &:= L(\psi_i^r, \psi_j^s) \text{ for } r, s > 1. \end{aligned}$$

The right-hand side \mathbf{g} is defined in a similar fashion: if $g_i^1 := G(\varphi_i^1)$ and $b_j^k := G(\psi_j^k)$ for $k = 1 \dots q-1$ and for $j = 1, \dots, N_k$,

$$\mathbf{g} = \begin{bmatrix} g^1 \\ b^1 \\ \vdots \\ b^{q-1} \end{bmatrix}$$

This problem also admits a unique solution, inheriting this property from the continuous formulation since $\mathcal{V}^q \subset H$. Solving the linear system in Eq. (5), one obtains the set of coefficients \mathbf{w} from which the approximate solution u^q is constructed through Eq. (4). This approach provides the same solution as if one had used a Galerkin approach at the finest resolution \mathcal{V}^q , just expressed in a different basis. The use of different resolutions brings, however, several advantages over the single-resolution system of equations: applications such as preconditioning and adaptive mesh

refinement can selectively adapt the number of wavelet coefficients to utilize. Yet, the matrix \mathbf{L} is difficult to deal with in practice. It is a large matrix with many non-zero off-diagonal elements, and with typically worsening condition numbers for large values of q . This issue renders the wavelet-Galerkin discretization not substantially better than the classical Galerkin discretization: in essence, many of the core properties of wavelets that make them extremely useful in signal processing are not as relevant in the FEM context. In other words, the performance of wavelet-based methods based on a multiresolution L^2 -orthogonal decomposition of the solution space [36, 37, 38] is affected by the regularity of wavelet coefficients because the decomposition is not inherently adapted to the underlying PDE.

2.2. Towards Operator Adapted Wavelets

Shortcomings of the wavelet-Galerkin approach have generated a large body of work aimed at alleviating the issues in solving Eq. (5) through the use of preconditioning, specific choices of wavelets to ease evaluation, small subsets/subbands of wavelets, etc. However, and despite a series of efforts to link wavelets to operators [5, 1], the foundations remained mostly unchanged in practical applications.

Operator-based multiresolution decomposition. In a sharp departure from the various palliative measures, Sudarshan [24] proposed to alter the nature of the multiresolution decomposition itself. He argued for what he called *scale-orthogonal wavelets*, i.e., FEM wavelets that should:

- (a). be operator-orthogonal to block-diagonalize the operator at play;
- (b). produce well-conditioned stiffness matrices (i.e., A^1 and B^k for $k = 1 \dots q-1$ in the previous example);
- (c). be localized or have fast decay for computational efficiency.

Property (a) implies a new multiresolution decomposition, no longer based on L^2 orthogonality, but on an operator-dependent notion of orthogonality: wavelets in \mathcal{W}^k should be \mathcal{L} -orthogonal to the basis functions of their own resolution and to all other wavelets. This implies that we now need to enforce:

$$\mathcal{V}^{k+1} = \mathcal{V}^k \oplus_{\mathcal{L}} \mathcal{W}^k,$$

where $\oplus_{\mathcal{L}}$ denotes direct sum of \mathcal{L} -orthogonal subspaces; the induced multiresolution decomposition becomes:

$$H \approx \mathcal{V}^q = \mathcal{V}^1 \oplus_{\mathcal{L}} \mathcal{W}^1 \oplus_{\mathcal{L}} \dots \mathcal{W}^{q-2} \oplus_{\mathcal{L}} \mathcal{W}^{q-1}$$

Numerical consequence of an operator-adapted decomposition. A wavelet-Galerkin treatment using such a decomposition would result in a significantly simpler linear solve: discretizing the weak problem in Eq. (2) now produces a linear system with the same structure as shown in Eq. (5), but the off-diagonal “inter-level” stiffness matrices all *vanish*, i.e. $M^{(r,s)} = \mathbf{0}$ for $r \neq s$. In other words, the matrix \mathbf{L} of the final linear system becomes *block-diagonal*. Consequently, the full linear system reduces to a set of q small independent linear equations $B^k w^k = b^k$ for $k = 1 \dots q-1$ and $A^1 v^1 = g^1$, each corresponding to its own resolution level. Property (b) additionally guarantees that each of these linear systems can be solved easily and reliably (Riesz stability). Finally, the last property (c) ensures that within a given resolution level, only nearby basis elements have non-zero contributions, rendering the matrices B^k sparse, and allowing for highly efficient storage as well as fast computations. As with classical wavelets, one obtains a frequency-based decomposition of the solution u^q , except that now this representation is truly adapted to the operator: the coarser basis functions derive from finer scales, changing their shapes accordingly to capture the eigenspaces of the corresponding operator as well as possible while staying spatially localized.

Wavelet construction. However, scale-orthogonal wavelets are not nearly as simple to construct as their use is desirable. A construction enforcing only (a) and (b) was proposed in [24]. A second construction was also offered in the same publication, this time by relaxing the exact \mathcal{L} -orthogonality of the decomposition, thus violating (a). However, this looser definition of wavelets leads to a construction involving Gram-Schmidt orthogonalizations of dense matrices, failing to deliver a fast algorithm except in restricted cases such as when the solution has very few non-zero wavelets coefficients. Recently, Owahdi [28] presented the first practical technique to achieve all three properties (a)-(b)-(c), to create operator adapted wavelets referred to as *gamblets* (due to their game theoretic interpretation). Gamblets have been used over the last year for alleviating the complexity bottleneck associated with parabolic and hyperbolic PDEs [29], dense kernel matrices [31], denoising PDEs [39] and eigenpairs computation [40].

2.3. Outline

In this paper, we revisit the recent work of Owhadi [28] and rederive his approach through an axiomatic derivation from a Finite Element Analysis point of view. We then demonstrate that this construction is general enough to handle functional spaces of differential forms based on Whitney basis functions [41, 42, 43] as well, thus offering a multiresolution analysis of arbitrary p -forms adapted to a given exterior operator. With this framework, one can construct novel computational tools such as vector-valued wavelets and divergence-free vector-valued basis functions adapted to any continuous, linear, bijective, and symmetric positive-definite operator. We also provide a detailed algorithmic description of our resulting construction of operator-adapted basis functions and wavelets of differential forms, and review how to use them for multiresolution analysis by leveraging the fact that each sub-band of the decomposition can be handled independently. Finally, we demonstrate the effectiveness of this fast computational procedure for numerical tasks such as coarse graining and model reduction of linear and non-linear partial differential equations.

3. Construction of operator-adapted wavelets

We present an algorithmic construction of operator-adapted wavelets through an axiomatic derivation formulated from a finite element perspective. The approach proceeds in a fine-to-coarse fashion, with locality and operator-orthogonality imposed directly based on simple linear algebra. For completeness as well as to offer a deeper insight into the optimality properties of our construction, we also provide in Appendix B a short summary of the alternative interpretation from [28] derived from a game-theoretical view of functional approximation and conditional generalized Gaussian processes.

3.1. Canonical multiresolution analysis

Our approach starts from an existing set of refinable basis functions whose spans form multiscale approximation spaces. Each basis function in this sequence of q nested approximation spaces is associated with an element of a mesh \mathcal{M}^k ($1 \leq k \leq q$), where meshes $\mathcal{M}^1, \mathcal{M}^2, \dots, \mathcal{M}^q$ provide increasingly finer spatial discretizations of the domain Ω .

Refinable basis functions. Let $\{\mathcal{V}^k\}_{k=1}^q \subset H^*$ be a nested sequence of q functional spaces, where each space \mathcal{V}^k is spanned by n_k compactly supported basis functions φ_i^k , i.e., $\mathcal{V}^k = \text{span} \{\varphi_i^k\}_{i=1}^{n_k}$. Each basis function φ_i^k is refinable for $1 \leq k < q$, that is to say, φ_i^k can be written as a weighted sum of finer basis functions $\{\varphi_j^{k+1}\}_{j=1}^{n_{k+1}}$:

$$\varphi_i^k = \sum_{j=1}^{n_{k+1}} \mathbf{C}_{ij}^k \varphi_j^{k+1} \quad (6)$$

Note that this relation between basis functions of two successive levels implies a specific topological refinement of the mesh \mathcal{M}^k at level k into a finer mesh \mathcal{M}^{k+1} . We do not make any special assumption on how these meshes are related (for instance, nested or not) as our construction is very general; as a consequence, we will rarely refer to these meshes, even if the definition of the basis functions are often expressed, explicitly or implicitly, in terms of barycentric coordinates over the elements of these meshes in practice. We do assume, however, that the refinement matrices \mathbf{C}^k defining the basis functions are *sparse*—and as we are about to see, we also require that there exists a sparse parametrization of their *kernels*. To distinguish these (user-specified) canonical basis functions from the operator-adapted basis functions we will construct at each level of resolution, we refer to $\{\varphi_i^k\}_{i=1}^{n_k}$ as *test* (or *measurement*) *functions* in the remainder of this paper: they will only be used to integrate against at the finest level to kickstart our construction of an operator-adapted multiresolution analysis.

Remark 1: Such a nested, or hierarchical, representation of the functional space \mathcal{V}^q is typical of Multiresolution Analysis (MRA) and wavelet-based adaptive solvers to allow for efficient handling of local refinements of the solution: our test functions are simply their ‘scaling functions’ (or ‘father wavelets’) — and any existing such function can be used in our construction. In practice, the finest functional space \mathcal{V}^q giving rise to the hierarchy is formed by the span of n_q linearly independent functions in L^2 . The L^2 -orthogonal complement of \mathcal{V}^k in \mathcal{V}^{k+1} is typically called \mathcal{W}^k , which is spanned by N_k wavelets $\{\psi_i^k\}_{i=1}^{N_k}$, with $N_k := n_{k+1} - n_k$ to capture the details at level k , i.e., the functions that are contained in \mathcal{V}^{k+1} , but not in \mathcal{V}^k . Note that only the test functions are needed in our construction, although their associated wavelets are implicitly present as well, as we will discuss when relevant.

Table 1: Summary of the main symbols used in our paper and their definitions.

Symbols	Meaning
\mathcal{V}^k	Space of test functions on resolution level k , $\dim \mathcal{V}^k = n_k$.
$\{\varphi_i^k\}_{i=1}^{n_k}$	Basis of the space of test functions \mathcal{V}^k .
\mathbf{C}^k	Refinement $n_k \times n_{k+1}$ matrix for test functions.
\mathcal{V}^k	Operator-adapted solution space on resolution level k , $\dim \mathcal{V}^k = n_k$. Note that $\mathcal{V}^q = \mathcal{V}^q$.
$\{\varphi_i^k\}_{i=1}^{n_k}$	Operator-adapted basis functions spanning solution space \mathcal{V}^k . Note that $\varphi_i^q = \varphi_i^q$.
\mathbb{C}^k	Refinement $n_k \times n_{k+1}$ matrix for operator-adapted basis functions, with $\varphi_i^k = \sum_{j=1}^{n_k} \mathbb{C}_{ij}^k \varphi_j^{k+1}$.
\mathcal{W}^k	Operator-adapted wavelet space on level k , with $\mathcal{V}^{k+1} = \mathcal{V}^k \oplus_{\mathcal{L}} \mathcal{W}^k$; $\dim \mathcal{W}^k = N_k := n_{k+1} - n_k$.
$\{\psi_i^k\}_{i=1}^{N_k}$	Operator-adapted wavelet basis spanning wavelet space \mathcal{W}^k .
\mathbf{W}^k	Refinement kernel $N_k \times n_{k+1}$ matrices with $\mathbf{C}^k \mathbf{W}^{k,T} = \mathbf{0}$, defining operator-adapted wavelets as $\psi_i^k = \sum_{j=1}^{N_k} \mathbf{W}_{ij}^k \varphi_j^{k+1}$.
\mathbf{A}^k	Stiffness matrix of size $n_k \times n_k$ for operator-adapted basis functions on level k , $\mathbf{A}_{ij}^k = L(\varphi_i^k, \varphi_j^k)$.
\mathbb{B}^k	Stiffness matrix of size $N_k \times N_k$ for operator-adapted wavelet basis on level k , $\mathbb{B}_{ij}^k = L(\psi_i^k, \psi_j^k)$.
u^q	Finite-element solution of $\mathcal{L}u = g$ (Eq.(1)) using finest solution space \mathcal{V}^q : $u^q = \sum_{i=1}^{n_1} v_i^1 \varphi_i^1 + \sum_{p=1}^{q-1} \sum_{i=1}^{N_p} w_i^p \psi_i^p$
u^k	Level- k approximation of u^q : $u^k \in \mathcal{V}^k$ and $u^k = \sum_{i=1}^{n_1} v_i^1 \varphi_i^1 + \sum_{p=1}^{k-1} \sum_{i=1}^{N_p} w_i^p \psi_i^p$
ω^k	Contribution to FEM solution from level- k wavelets: $\omega^k \in \mathcal{W}^k$ and $\omega^k = u^{k+1} - u^k = \sum_{i=1}^{N_k} w_i^k \psi_i^k$

Refinement kernel. From a refinement matrix \mathbf{C}^k , we define a sparse, $N_k \times n_{k+1}$ *refinement kernel matrix* \mathbf{W}^k satisfying:

$$\mathbf{C}^k \mathbf{W}^{k,T} = \mathbf{0}_{n_k \times N_k}.$$

In other words, the N_k rows of \mathbf{W}^k form a basis for $\text{Ker } \mathbf{C}^k$, hence our naming convention. This matrix should only have a few non-zero entries per row, corresponding to nearby test functions in the physical domain. In addition, the condition number of \mathbf{W}^k should be as close to 1 as possible to achieve lower condition numbers for the operator-adapted stiffness matrices we will construct (corresponding to the B^k matrices in the previous section).

Remark 2: If the test (scaling) functions φ_i^k were L^2 -orthonormal, the matrix \mathbf{W}^k would correspond to the refinement relation of the *wavelets* associated to the scaling functions, i.e.,

$$\psi_i^k = \sum_{j=1}^{n_{k+1}} \mathbf{W}_{ij}^k \varphi_j^{k+1}.$$

Indeed, for a typical pair of mother and father wavelets, one has: $\langle \varphi_i^k, \psi_j^k \rangle_{L^2} = \sum_{r,s} \mathbf{C}_{ir}^k \langle \varphi_r^{k+1}, \varphi_s^{k+1} \rangle_{L^2} \mathbf{W}_{sj}^{k,T}$. So if the mass matrix at level $k+1$ of the test functions is the identity, the wavelets form precisely the L^2 -orthogonal subspace \mathcal{W}^k mentioned in Rmk. 1. However, our construction does *not* require L^2 -orthonormality of the test functions, so the refinement kernel matrices \mathbf{W}^k do not directly define L^2 -orthogonal wavelets: these matrices will be used instead to *parameterize the kernel of* \mathbf{C}^k in order to construct our operator-adapted basis functions and wavelets.

3.2. Setup for operator-adapted basis functions and wavelets

From the refinement equations of the user-selected test functions and a choice of refinement kernel matrices, we wish to bootstrap a fine-to-coarse construction of operator-adapted basis functions and associated wavelets. We begin by listing a minimal set of basic, desirable axiomatic properties that are formulated to enforce that the new basis functions and wavelets define an operator-adapted multiresolution analysis of the solution space, and are easy to construct via linear algebra.

Cardinality of degrees of freedom. Our first requirement is that our operator-adapted multiscale decomposition must share the *same* cardinality as the canonical multiresolution analysis based on the test functions φ_i^k . Consequently, for each level $k = 1 \dots q$, we need to construct n_k operator-adapted basis functions, which we will refer to as $\{\varphi_i^k\}_{i=1}^{n_k}$, spanning a functional space that we will denote \mathcal{V}^k .

Collocation of degrees of freedom. One of the key properties of MRA is spatial localization: unlike operator eigenfunctions that are not attached to any particular spatial location, basis functions are always associated with a particular element (which, consequently, makes the wavelets localized as well). We exploit the existence of the spatially-localized canonical test functions φ_i^k by imposing a weak collocation of the operator-adapted basis functions (we use

δ_{ij} to denote the Kronecker delta):

$$\langle \varphi_i^k, \varphi_j^k \rangle_{L^2} = \delta_{ij} \quad \text{for all } i = 1 \dots n_k, j = 1 \dots n_k, k = 1 \dots q, \quad (7)$$

which will ensure that our operator-adapted basis functions φ_i^k are similarly located in the domain.

Operator-adapted refinability. The new basis functions φ_i^k should be refinable to induce a hierarchy of nested functional spaces. We denote by \mathbb{C}^k the corresponding (unknown) refinement matrices in order to impose the following refinement relation between levels:

$$\varphi_i^k = \sum_{j=1}^{n_{k+1}} \mathbb{C}_{ij}^k \varphi_j^{k+1}. \quad (8)$$

Operator-adapted wavelets. We wish to construct wavelets as well, associated with the operator-adapted basis functions. Let $\{\psi_i^k\}_{i=1}^{N_k}$ be the $N_k := n_{k+1} - n_k$ (pre-)wavelets defined at each level k , and let \mathbb{W}^k the corresponding functional spaces they span. These wavelets should be \mathcal{L} -orthogonal to the operator-adapted basis functions of the same level, to enforce that $\mathbb{V}^{k+1} = \mathbb{V}^k \oplus_{\mathcal{L}} \mathbb{W}^k$; consequently, one must have:

$$\langle \varphi_i^k, \psi_j^k \rangle_{\mathcal{L}} = 0 \quad \text{for all } i = 1 \dots n_k, j = 1 \dots N_k, k = 1 \dots q - 1. \quad (9)$$

Wavelet refinability. Since we assumed that $\mathbb{W}^k \subset \mathbb{V}^{k+1}$, our adapted wavelets at level k must also be a linear combinations of adapted basis functions at level $k+1$. We require that the corresponding wavelet refinement matrices are, in fact, the *refinement kernel matrices* \mathbb{W}^k , that is,

$$\psi_i^k = \sum_{j=1}^{n_{k+1}} \mathbb{W}_{ij}^k \varphi_j^{k+1}. \quad (10)$$

This particular choice of refinement matrix for our wavelets allows us to leverage the kernel of the canonical refinement from which we derive the operator-adapted decomposition.

Stiffness matrices. By analogy to the L^2 -orthogonal case described earlier, we denote by \mathbb{A}^k the stiffness matrix of our new basis functions, and by \mathbb{B}^k the stiffness matrix of the new wavelets. That is,

$$\mathbb{A}_{ij}^k = L(\varphi_i^k, \varphi_j^k), \quad \mathbb{B}_{ij}^k = L(\psi_i^k, \psi_j^k). \quad (11)$$

Note finally that our notational conventions were selected such that bold letters denote user-selected entities, while outlined letters refer to entities we must construct. In particular, the operator-adapted refinement matrices \mathbb{C}^k are the key ingredients: computing them allows for the construction of adapted basis functions, which in turn leads to the construction of their adapted wavelets, and thus of the induced stiffness matrices; see Table 1 for a review of the various symbols and their definitions.

3.3. Resulting linear algebraic conditions

From the axiomatic characterization of our \mathcal{L} -adapted multiresolution analysis, we can easily derive linear algebra conditions on the matrices involved in our construction so that Properties (a)-(c) from Section 2.2 are satisfied. Indeed,

- the collocation conditions expressed in Eq. (7) directly implies that

$$\mathbb{C}^k \mathbf{C}^{k,T} = \mathbb{I}_{n_k}; \quad (12)$$

- due to Eq. (10), the \mathcal{L} -orthogonality conditions expressed in Eq. (9) can be rewritten as:

$$\mathbb{C}^k \mathbb{A}^{k+1} \mathbf{W}^{k,T} = \mathbf{0}_{n_k \times N_k}; \quad (13)$$

- finally, the stiffness matrices defined in Eq. (11) are linked through \mathbf{W}^k due to Eq. (10):

$$\mathbb{B}^k = \mathbf{W}^k \mathbb{A}^{k+1} \mathbf{W}^{k,T}. \quad (14)$$

From these conditions, we can establish a closed-form expression for \mathbb{C}^k satisfying all of our requirements:

Proposition: For Eqs. (12), (13), and (14) to hold, one must have:

$$\mathbb{C}^k = \mathbf{C}^{k,\dagger} \left[\mathbb{I}_{n_{k+1}} - \mathbb{A}^{k+1} \mathbf{W}^{k,T} \left(\mathbb{B}^k \right)^{-1} \mathbf{W}^k \right], \quad \text{where } \mathbf{C}^{k,\dagger} = (\mathbf{C}^k \mathbf{C}^{k,T})^{-1} \mathbf{C}^k. \quad (15)$$

Proof: Given that $\mathbb{C}^k \mathbf{C}^{k,T} = \mathbb{I}_{n_k}$ (Eq. (12)), the refinement matrices \mathbb{C}^k must be of the form

$$\mathbb{C}^k = \mathbf{C}^{k,\dagger} + \mathbb{Z} \mathbf{W}^k,$$

where $\mathbf{C}^{k,\dagger}$ denotes the pseudoinverse of \mathbf{C}^k and \mathbb{Z} is an arbitrary $n_k \times N_k$ matrix: indeed, we saw in Section 3.1 that the rows of \mathbf{W}^k form a basis for $\text{Ker } \mathbf{C}^k$ so that $\mathbf{W}^k \mathbf{C}^{k,T} = \mathbf{0}$. Additionally, we must have $\mathbb{C}^k \mathbb{A}^{k+1} \mathbf{W}^{k,T} = \mathbf{0}$ (Eq. (13)), implying that

$$(\mathbf{C}^{k,\dagger} + \mathbb{Z} \mathbf{W}^k) \mathbb{A}^{k+1} \mathbf{W}^{k,T} = \mathbf{0} \quad \implies \quad \mathbb{Z} = -\mathbf{C}^{k,\dagger} \mathbb{A}^{k+1} \mathbf{W}^{k,T} \left(\mathbf{W}^k \mathbb{A}^{k+1} \mathbf{W}^{k,T} \right)^{-1}.$$

Substituting \mathbb{B}^k for $\mathbf{W}^k \mathbb{A}^{k+1} \mathbf{W}^{k,T}$ (Eq. (14)) and plugging \mathbb{Z} back into the expression of \mathbb{C}^k leads to the proposition. \square

3.4. Computational procedure for our operator-adapted wavelet-Galerkin approach

Based on the closed-form expression of the refinement matrices \mathbb{C}^k , we can now derive the computational procedure to turn the weak form in Eq. (5) on the finest level \mathcal{V}^q into q independent and well-conditioned linear systems through our \mathcal{L} -adapted MRA.

Bootstrapping the finest scale. In order to start our construction, we need to select the basis functions $\{\varphi_i^q\}_{i=1}^{n_q}$ at the finest level. We directly pick them to be the finest test functions *unaltered*, i.e., we set

$$\varphi_i^q = \varphi_i^q \quad (16)$$

for simplicity, as it trivially enforces that these initial fine basis functions are attached to the same mesh elements as their canonical equivalents. Note that this choice implicitly amounts to rescaling the test functions at the finest scale through the inverse of the mass matrix, which is a valid alternative as it formally enforces Eq. (7); one can also see this choice as a simple way to ensure that the finest basis functions are collocated with the test functions without actually having to enforce Eq. (7). The resulting stiffness matrix \mathbb{A}^q is then trivially evaluated as it is precisely the usual finite-element stiffness matrix \mathbf{A}^q , and the right hand side of the original problem is represented as a vector g^q by integrating g against these finest basis functions, i.e., $g_i^q = G(\varphi_i^q)$.

Fine-to-coarse evaluation. Given our setup and the closed-form expression of the key refinement matrix, we can now proceed hierarchically, in a bottom-up fashion, by evaluating relevant matrices and vectors in the following order:

$$\left\{ \begin{array}{l} \mathbb{B}^{k-1} = \mathbf{W}^{k-1} \mathbb{A}^k \mathbf{W}^{k-1,T} \\ \mathbb{C}^{k-1} = \mathbf{C}^{k-1,\dagger} \left[\mathbb{I}_{n_k} - \mathbb{A}^k \mathbf{W}^{k-1,T} \left(\mathbb{B}^{k-1} \right)^{-1} \mathbf{W}^{k-1} \right] \\ \mathbb{A}^{k-1} = \mathbb{C}^{k-1} \mathbb{A}^k \mathbb{C}^{k-1,T} \\ b^{k-1} = \mathbf{W}^{k-1} g^k \\ g^{k-1} = \mathbb{C}^{k-1} g^k. \end{array} \right. \quad (17)$$

Resulting weak form. Once these evaluations have been done all the way to level 1, the weak form in Eq. (5) is, by construction, equivalent to q independent (and, as detailed later, well-conditioned) linear systems:

$$\left\{ \begin{array}{l} \mathbb{A}^1 v^1 = g^1 \\ \mathbb{B}^k w^k = b^k \quad \text{for } k = 1 \dots q-1 \end{array} \right. \quad (18)$$

from which the *level-k FEM solution* $u^k \in \mathcal{V}^k$ for $k = 2 \dots q$ can then be evaluated via

$$u^k = \sum_{i=1}^{n_1} v_i^1 \varphi_i^1 + \sum_{p=1}^{k-1} \sum_{i=1}^{N_p} w_i^p \psi_i^p; \quad (19)$$

compare to the large linear system in Eq. (5) using the usual L^2 -orthogonal MRA. In particular, due to the independence of the linear systems in Eq. (18), one can easily compute approximations of the fine solution u^q by solving only the first k linear systems (instead of the full q linear systems) and returning u^k using Eq. (19).

3.5. Properties of the construction

As mentioned earlier, the fine-to-coarse construction we derived above from basic desirable axioms is, in fact, equivalent to the work of [28], originally developed from a game-theoretical approach to optimal functional approximation. They argued that the optimal operator-adapted basis functions should satisfy a variational definition, formulated as:

$$\varphi_i^k = \underset{\phi \in H}{\operatorname{argmin}} \|\phi\|_{\mathcal{L}}^2 \text{ s. t. } \langle \phi, \varphi_j^k \rangle_{L^2} = \delta_{ij} \text{ for } j = 1 \dots n_k \quad (20)$$

This variational formulation implies that the basis functions φ_i^k can also be identified as optimal recovery splines in the sense of [44], and viewed as a generalization of energy minimizing splines in the sense of [45], variational multiscale (or LOD) basis functions in the sense of [46, 47] and polyharmonic splines in the sense of [48, 49, 50]. For completeness we provide a short summary of their arguments leading to this definition in Appendix B. From this variational definition, they derive a construction for operator-adapted basis functions and wavelets that mirrors ours, reinforcing the fact that our axiomatic characterizations are quite natural. As a consequence, we inherit the properties of the functions resulting from Eq. (20) as well. These properties include in particular:

- the condition numbers of stiffness matrices \mathbb{B}^k are uniformly bounded for all k ;
- the resulting basis functions φ_i^k decay exponentially fast under reasonably mild assumptions on the interactions between test functions and operator;
- operator-adapted basis functions allow for small approximation errors even if high frequencies (i.e., fine levels) are ignored, confirming their coarse-graining properties.

We provide a summary of the analysis leading to these conditions in Appendix B too, in order to make the underlying conditions and the actual properties more explicit, but interested readers should refer to the original paper [28] and to [32] for rigorous proofs. Note that we will numerically verify these properties in our context, and will explain the practical implications of their assumptions.

3.6. Practical implementation

While the fine-to-coarse construction described in Sec. 3.4 is simple, it fails to be efficient to implement *as is* since many of the matrices involved are not sparse in general. Yet, the fact that operator-adapted basis functions are decaying exponentially fast can be leveraged to enforce sparsity throughout the construction. Additionally, we did not explicitly describe how the \mathcal{L} -adapted basis functions and wavelets are constructed in practice to recover the solution from Eq. (19), so we now go over implementation details of the fine-to-coarse construction procedure.

Algorithm 1 Operator-Adapted Basis and Wavelet Construction

Input: Basis matrix \mathbb{D}^q and associated stiffness matrix \mathbb{A}^q on finest resolution level, refinement matrices $\{\mathbf{C}^k, \mathbf{W}^k\}_{k=1}^{q-1}$.

for $k = q$ to **2 do**

$\mathbb{D}^{k-1}, \Psi^{k-1}, \mathbb{A}^{k-1}, \mathbb{B}^{k-1} \leftarrow$ Per-level Coarsening Step ($\mathbb{D}^k, \mathbb{A}^k, \mathbf{C}^{k-1}, \mathbf{W}^{k-1}$)

end for

Output: Operator-adapted bases $\{\mathbb{D}^k\}_{k=1}^q$, pre-wavelets $\{\Psi^k\}_{k=1}^{q-1}$, and their stiffness matrices $\{\mathbb{A}^k\}_{k=1}^q$ and $\{\mathbb{B}^k\}_{k=1}^{q-1}$.

Basis functions and wavelets. Once the finest \mathcal{L} -adapted basis functions $\{\varphi_i^q\}_{i=1}^{n_q}$ have been defined (Eq. (16)), it is clear from Eq. (8) and Eq. (10) that all other basis functions and wavelets are simply *linear combinations* of these fine functions. Therefore, we can store them as just the coefficients of the fine basis functions they are made of. To simplify the notation, we assemble a (sparse) vector \mathbb{D}^k of operator-adapted basis functions (with n_k rows) and a (sparse) vector Ψ^k of pre-wavelets (with N_k rows) for each resolution level k to keep track of these linear combinations:

$$\mathbb{D}^k = [\varphi_1^k \ \varphi_2^k \ \dots \ \varphi_{n_k}^k]^T, \quad \Psi^k = [\psi_1^k \ \psi_2^k \ \dots \ \psi_{N_k}^k]^T,$$

where each φ_i^k and ψ_j^k is discretized as coefficients associated to all the n_q basis functions of the fine mesh. During the fine-to-coarse construction, these vectors are trivially assembled starting from the finest level basis $\{\varphi_i^q\}_{i=1}^{n_q}$ and onto coarser resolutions via the refinement relations of Eqs. (8) and (10). The pseudocode for computing adapted basis functions and corresponding wavelets is given in Alg. 1.

Algorithm 2 Per-level Coarsening Step

Input: Basis matrix Φ^k and its stiffness matrix \mathbb{A}^k on k -th resolution level, refinement matrices \mathbf{C}^{k-1} , \mathbf{W}^{k-1} .

Compute wavelet stiffness: $\mathbb{B}^{k-1} \leftarrow \mathbf{W}^{k-1} \mathbb{A}^k \mathbf{W}^{k-1,T}$

Compute pseudoinverse $\mathbf{C}^{k-1,\dagger} \leftarrow (\mathbf{C}^{k-1} \mathbf{C}^{k-1,T})^{-1} \mathbf{C}^{k-1}$

Compute intermediate matrix \mathbb{Z} through $\mathbb{B}^{k-1} \mathbb{Z}^T = -\mathbf{W}^{k-1} \mathbb{A}^k \mathbf{C}^{k-1,\dagger T}$ (fast solve through localization)

Compute adapted basis refinement matrix: $\mathbf{C}^{k-1} \leftarrow \mathbf{C}^{k-1,\dagger} + \mathbb{Z} \mathbf{W}^{k-1}$

Compute adapted wavelets on level $(k-1)$: $\Psi^{k-1} \leftarrow \mathbf{W}^{k-1} \Phi^k$

Compute adapted basis on level $(k-1)$: $\Phi^{k-1} \leftarrow \mathbf{C}^{k-1} \Phi^k$

Assemble stiffness matrix for lower resolution level: $\mathbb{A}^{k-1} \leftarrow \mathbf{C}^{k-1} \mathbb{A}^k \mathbf{C}^{k-1,T}$ (sparsified through truncation)

Output: Operator-adapted basis matrix Φ^{k-1} , pre-wavelets Ψ^{k-1} , stiffness matrices \mathbb{A}^{k-1} , \mathbb{B}^{k-1} and refinement \mathbf{C}^{k-1} .

Sparsification. As a consequence of the exponential decay of operator-adapted basis functions discussed in Sec. 3.5, we can reduce computational complexity without sacrificing numerical precision by working with sparse matrices throughout the fine-to-coarse construction. Thresholding near-zero values will have no significant numerical effects, and will allow for efficient sparse computations; we refer the reader to [28, 32] for a thorough discussion of accuracy and error propagation across scales induced by this procedure. In the coarsening algorithm from level k to level $k-1$, we thus sparsify both matrix \mathbf{C}^{k-1} and matrix \mathbb{A}^{k-1} through simple truncation.

Fast evaluation of \mathbf{C}^{k-1} . Seemingly, the most tedious part of the coarsening procedure is the computation of the \mathcal{L} -adapted refinement matrix \mathbf{C}^k from Eq. (15). Evaluation is done in two stages: one linear solve to derive an intermediate matrix \mathbb{Z} (used in the proof of the Proposition in Sec. 3.3), then the final expression based on \mathbb{Z} , $\mathbf{C}^{k-1,\dagger}$, \mathbb{A}^k and \mathbf{W}^{k-1} , see Alg. 2. In [28], \mathbb{Z} is evaluated efficiently through localization, and we adopt the same approach here: in order to efficiently solve the linear system $\mathbb{B}^{k-1} \mathbb{Z}^T = -\mathbf{W}^{k-1} \mathbb{A}^k \mathbf{C}^{k-1,\dagger T}$, we localize both \mathbb{B}^{k-1} and the right hand side by exploiting the sparsity of \mathbf{W}^{k-1} . We compute each i -th column \mathbb{Z}_i of \mathbb{Z}^T by solving a small linear system as follows. Since the column \mathbb{Z}_i corresponds to the i -th element of the coarser mesh, consider all mesh elements that are less than three element away¹, inducing a small region R of the domain. We first assemble a set of indices S such that $\ell \in S$ iff the ℓ -th row of \mathbf{W}^{k-1} has at least one non-zero coefficient on one of the finer mesh elements in R . We then assemble a matrix \mathbb{B} which is a submatrix of the stiffness \mathbb{B}^{k-1} using only the rows and columns whose indices are in S (this is a reduced stiffness matrix of the wavelets around i whose support intersects R). We also assemble a vector \mathbf{z}_i by subsampling the i -th column of $\mathbf{W}^{k-1} \mathbb{A}^k \mathbf{C}^{k-1,\dagger T}$ using the indices from S . Now entries of \mathbb{Z}_i with indices from S , assembled in a vector \mathbf{z}_i , are evaluated by solving the small linear system: $\mathbb{B} \mathbf{z}_i = -\mathbf{z}_i$, while all other entries of \mathbb{Z}_i are set to zero. This procedure to compute \mathbb{Z} is significantly faster than a full-blown SparseLU or Preconditioned Conjugate Gradient solve as it only involves linear solves of small size, and less prone to inaccuracies than a truncation-based sparsification of the various matrices involved. In addition, it promotes sparsity for \mathbb{Z} and of all subsequent operations, e.g., for the refinement matrix \mathbf{C}^{k-1} and the basis functions stored as Φ^{k-1} .

Complexity. Exploiting sparsity makes the construction extremely efficient: in dimension d , it has computational complexity of $\mathcal{O}(n_q \log^{2d+1} n_q)$ [28, 32] if one uses a fine grid with n_q degrees of freedom. Our timing tests confirm this expected complexity of our algorithm with fast evaluation of \mathbf{C}^k . However, we observed that our python implementation can produce slightly worse timings for $n_q < 10^3$ compared to regular sparse solvers combined with naive threshold-based sparsification: our fast evaluation only pays off for large enough problems.

3.7. Discussion

Pseudocode of the numerical procedure to compute an operator-adapted decomposition is given in Algs. 1 and 2. Notice that since the operator-adapted basis functions and wavelets are stored as linear combinations of the original test functions on the finest level \mathcal{M}^q , their evaluation at any point of the domain is fast as it only involves a few local evaluations of test functions. Consequently, all the regular finite-element approaches used to solve partial differential equations can be applied at any level of approximation—except that our “operator-aware” basis functions are used in lieu of the usual polynomial basis functions, just like in other works which derive tailored shape functions to improve

¹Note that this spatial extent can be adapted on a per-level basis to tailor the sparsity vs. accuracy balance of the decomposition.

Algorithm 3 Linear Solve

Input: Basis matrix Φ^q and its stiffness matrix \mathbb{A}^q on finest resolution level, refinement matrices $\{\mathbf{C}^k, \mathbf{W}^k\}_{k=1}^{q-1}$, right-hand side g^q in basis Φ^q .

for $k = q$ **to** 2 **do**

$\Phi^{k-1}, \Psi^{k-1}, \mathbb{A}^{k-1}, \mathbb{B}^{k-1}, \mathbb{C}^{k-1} \leftarrow$ Per-level Coarsening step ($\Phi^k, \mathbb{A}^k, \mathbf{C}^{k-1}, \mathbf{W}^{k-1}$)

Find wavelet coefficients w^{k-1} by solving $\mathbb{B}^{k-1} w^{k-1} = \mathbf{W}^{k-1} g^k$

Assemble level $(k-1)$ wavelet solution $\omega^{k-1} \leftarrow \Psi^{k-1, T} w^{k-1}$

Coarsen right hand side $g^{k-1} \leftarrow \mathbb{C}^{k-1} g^k$

end for

Final coefficient solve $\mathbb{A}^1 v^1 = g^1$

Assemble final solution $u^q \leftarrow \Phi^{1, T} v^1 + \omega^1 + \dots + \omega^{q-1}$

Output: Solution to finite element problem u^q .

Algorithm 4 Nonlinear Solve

Input: Basis matrix Φ^q , refinement matrices $\{\mathbf{C}^k, \mathbf{W}^k\}_{k=1}^{q-1}$, stopping time $MaxIter$ and basis recomputation rate m .

for $s = 0$ **to** $MaxIter$ **do**

Compute stiffness \mathbb{A}^q of linearized operator at current time step in basis Φ^q

Compute right hand side g^q at current time step in basis Φ^q

if $(s \bmod m == 0)$ **then**

Compute operator-adapted bases/wavelets $\Phi^1, \{\Psi^k\}_{k=1}^{q-1}$ and stiffness matrices $\mathbb{A}^1, \{\mathbb{B}^k\}_{k=1}^{q-1}$.

Optional: perform model reduction by selecting a subset of operator-adapted wavelets as active.

end if

Find solution $u^q(s)$ at current time step using most recent active operator-adapted wavelets.

$s \leftarrow s + 1$

end for

Output: Solution to finite element problem $u^q(s)$.

accuracy on coarse grids [51]. The whole decomposition (or, for a fast approximation, part of the decomposition) can also be used to solve a linear differential equation like Eq. (1) by adapting the wavelet-Galerkin approach to our framework, as explicitly described in Alg. 3. Finally, non-linear differential equations such as $\mathcal{L}(u) = g$ can also be solved efficiently with our approach: by linearizing the non-linear operator at the current estimate of u and using the associated adapted basis functions and wavelets, one can iteratively improve the approximation of the solution as described in Alg. 4.

While we implicitly assumed that our test functions are traditional node-based and scalar-valued functions, we next describe how all the expressions given thus far actually apply to a much larger class of finite elements: our decomposition applies to differential forms and linear self-adjoint positive-definite exterior operators.

4. Operator-adapted Wavelets for p -Forms

Finite element formulations have most commonly used node-based basis functions to discretize both scalar functions and the coordinates of vector fields. A growing trend in recent years consists in employing mixed finite element methods, where basis functions are associated with arbitrary mesh elements (vertices, edges, faces, and cells for 3D grids). These methods have been instrumental in promoting the use of finite dimensional spaces of *differential forms* for coordinate-free computations with scalar and vector fields. We demonstrate in this section that our wavelet construction extends to the *multiresolution analysis of differential p -forms adapted to a given continuous, linear, bijective, and self-adjoint positive-definite exterior operator* (see three examples in Fig. 1). We also exhibit basis functions and wavelets of p -forms adapted to common operators.

4.1. Finite element differential forms

In Cartan's "Exterior Calculus" [33], differential forms are used as the building blocks of differential and integral calculus in arbitrary dimension. For computational purposes, a number of efforts have been dedicated to derive a

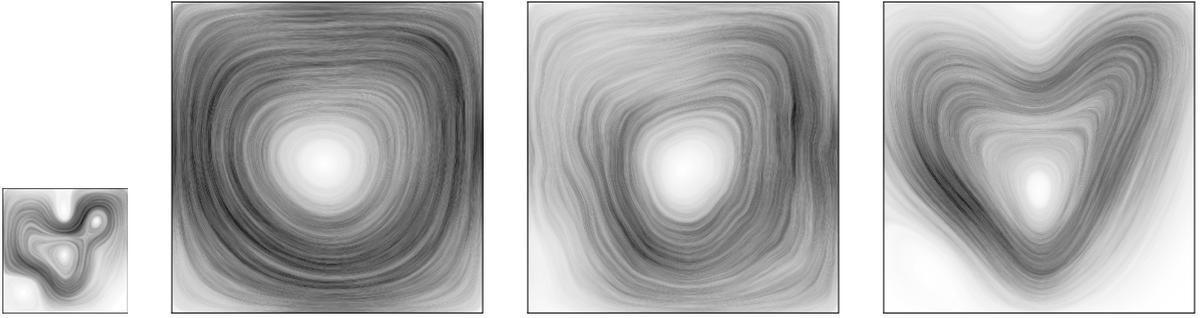


Fig. 1: **Characteristic Solutions.** Elements of the solution spaces of the 1-form Laplacian (left), the modified 1-form Laplacian (middle, see Eq. (28)) and the advection-diffusion (right, see vector field in Fig. 19(top)) operators, computed by solving $\mathcal{L}u = g$ for the same right-hand side 1-form g (depicted far left).

discrete version of this coordinate-free calculus through the use of finite dimensional approximation of differential forms.

Differential forms for computations. Mixed finite elements were introduced in the mid-1970s as an effort to build mixed discrete formulations for second order elliptic operators [52, 53]. The realization that they, in fact, were finite element spaces of differential forms came only later in [54]. At the root of these discrete differential forms is algebraic topology: cochains (resp., chains) provide a natural discretization of differential forms (resp., domain of integration) [55, 56]. Since a set of scalar values on vertices, edges, faces, and cells are proper discrete counterparts of respectively pointwise function evaluations, line integrals, surface integrals, and volume integrals [41], one can emulate exterior calculus on finite grids through the construction of finite-element differential forms [57]. This point of view naturally encompasses the use of volume integrals in finite volume methods and scalar functions in finite element methods, and also includes non-conforming “edge elements” and “facet elements” which are conforming in the spaces H_{div} and H_{curl} [53, 58]. Generalized basis functions that span the spaces of differential forms are particularly attractive as many important scalar and vector calculus properties and invariants of the continuous setting directly carry over to the discrete world in a coordinate-free manner, culminating in a discrete Hodge theory [43, 57, 34]. The resulting finite-element framework for differential forms provides a powerful tool for solving a wide range of partial differential equations, which has been proven useful in a variety of applications such as computational electromagnetism [41, 59], fluid simulation [53, 60, 61, 62, 63], quadrangle meshing of curved manifolds [64], and computer vision [65] to mention a few.

Whitney forms as basis functions of differential forms. The simplest basis of finite-element differential forms happen to correspond to a construct in algebraic topology known as Whitney forms [54]. Their initial purpose was to relate the de Rham sequence of smooth differential forms to simplicial cochain sequences in order to prove de Rham’s theorem, which states that these sequences have isomorphic cohomology groups [66, 67]. But in a computational setting, they can be seen as low-order “basis functions” that turn values on mesh elements into differential forms: given a vector $U = (u_1, \dots, u_n)^T$ of real values on all the n p -dimensional elements $\{\sigma_i\}_{i=1}^n$ of a mesh, one can construct a p -form \mathbf{u} as a linear combination of the Whitney forms $(\phi_1^{[p]}, \dots, \phi_n^{[p]})$ associated with these mesh elements through

$$\forall x \in \Omega, \quad \mathbf{u}(x) = \sum_{i=1}^n u_i \phi_i^{[p]}(x) = \Phi^{[p]} U, \quad (21)$$

where $\Phi^{[p]}$ stores all the Whitney p -forms as a row vector. In other words, there is one Whitney p -form for each p -dimensional element of a given mesh \mathcal{M} , and their span generates a finite-dimensional space of differential p -forms. A Whitney p -form $\phi_i^{[p]}$ is expressed in terms of the barycentric coordinates of the associated p -element σ_i , and its integration over this (oriented) p -element is

$$\int_{\sigma_j} \phi_i^{[p]} = \delta_{ij}, \quad (22)$$

making it “histopolating”—an extension of the notion of pointwise interpolation where integrals of a function over given regions match a prescribed set of values, see, e.g., [68, 69, 70]. When a metric is introduced on the ambient space, Whitney forms can be expressed as either scalar-valued (for $p = 0$ and 3 in \mathbb{R}^3) or, via musical isomorphisms (♭ and ‡, see [71]), vector-valued (for $p = 1$ and 2 in \mathbb{R}^3) functions with local support.

From low-order to high-order Whitney forms. While Whitney forms provide low-order interpolation of cochains (in particular, 0-forms are piecewise linear on a triangulation and trilinear on a cubical grid), a number of higher-order variants of these bases of differential forms have since then been proposed over simplicial or polytopal meshes [72, 73, 74, 75, 34, 76, 77, 78, 79, 80], some even with spectral accuracy [81, 82]. Their constructions vary greatly: a few are derived from tensor products of Whitney forms, while others follow the traditional approach of using higher “moments” to define various degrees of freedom associated to mesh elements in order to raise the order of the bases of differential forms while enforcing proper continuity across elements. Such high-order basis functions are carefully designed to span finite-dimensional differential p -form spaces Λ_p that form a *de Rham complex*,

$$\Lambda_0(\mathcal{M}) \xrightarrow{(\nabla)} \Lambda_1(\mathcal{M}) \xrightarrow{(\nabla \times)} \Lambda_2(\mathcal{M}) \xrightarrow{(\nabla \cdot)} \Lambda_3(\mathcal{M}) \rightarrow \mathbf{0}, \tag{23}$$

in which the fundamental differential operators gradient, curl, and divergence are unified as the exterior derivative \mathbf{d} (which satisfies $\mathbf{d} \circ \mathbf{d} = \mathbf{0}$), to offer structure-preserving finite element computations.

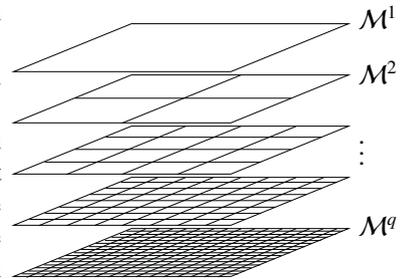
4.2. Refinable test functions for differential forms

The notion of refinability of basis functions [83] is crucial to the generation of wavelets, and even our multiresolution analysis requires an existing set of refinable test functions to bootstrap the bottom-up construction of our operator-adapted wavelets. Alas, refinability is a property rarely sought after or even discussed in the context of basis functions of differential p -forms: while the original Whitney forms are in fact refinable as pointed out in [75], most higher-order extensions are not known to have this property. As a consequence, the construction of wavelets for differential forms remains a challenge despite obvious applications for efficient numerical computations in electromagnetism and other computational fields. To the authors’ knowledge, the only arbitrary-order Whitney bases that are refinable by construction on both simplicial and polytopal meshes are the subdivision-based forms defined in [75] and extended in [84]. We thus review their construction, usage, and properties as this particular family of refinable basis functions of differential forms and their associated mesh hierarchy will be used as test functions for differential forms in our multiresolution analysis.

4.2.1. Mesh hierarchy

Due to the large body of known subdivision schemes for simplicial or polytopal meshes of arbitrary manifold domains [85], any (nested or non-nested) hierarchy of meshes for which subdivision schemes exist could be used for our construction. However, regular (Cartesian) grids are undeniably simpler due to their tensor product nature, and often more amenable to efficient implementation. Consequently, for simplicity of presentation, we will only consider Cartesian meshes of simply connected domains in our explanations.

We thus assume that the computational domain is a d -dimensional cube in \mathbb{R}^d (for $d = 1, 2$, or 3), over which a nested hierarchy of regular grids $\{\mathcal{M}^k\}_{k=1}^q$ is defined, with \mathcal{M}^{k+1} being twice as fine as \mathcal{M}^k such that each p -dimensional mesh element of \mathcal{M}^k is the union of 2^p p -dimensional mesh element of \mathcal{M}^{k+1} for $p = 1, 2$, or 3 (see inset for a 2D example of such a hierarchy). Each mesh element is given an arbitrary orientation, so that the boundary of any element is a linear combination of its faces with coefficients ± 1 based on whether the relative orientations match [86]. Considering only a simple d -dimensional cube domain is not a very restrictive assumption: we will show that our construction can not only homogenize operators, but also the spatial domain over which they are applied, making this simpler case sufficient to deal with complex domains (see Figs. 8 & 9 and Sec. 4.5)—but for accurate boundary handling, meshes conforming to curved boundaries (e.g., [87]) are obviously preferable.



4.2.2. (Low-order) Whitney forms ϕ as test functions

Whitney bases for p -forms on cubical grids have been known for decades [41]: Whitney 0-forms $\phi^{[0]}$ associated to vertex values are the standard trilinear functions over each hexahedral element, Whitney 1-forms $\phi^{[1]}$ associated with integrals over edges span the H_{curl} -conforming space, while Whitney 2-forms $\phi^{[2]}$ associated with integrals over faces span the H_{div} -conforming space. Finally, Whitney 3-forms $\phi^{[3]}$ associated with integrals over cells are equal to the indicator function of their associated cell. As noted in [88], these low-order Whitney forms are in fact refinable. Therefore, the Whitney forms $\phi^{[p]}$ for $p=0, \dots, 3$ can be used as test functions at any level of the mesh hierarchy, with one Whitney test function $\phi_i^{[p]}$ for each i -th p -element: the finite-dimensional differential p -form space Λ_p plays the role of the functional space of test functions at any level.

4.2.3. Design of higher-order Whitney forms through subdivision

From these low-order original Whitney forms, higher-order refinable Whitney forms can be constructed as well using the concept of “subdivision” [88]. A subdivision scheme for p -forms maps scalar coefficients on p -elements from a coarse mesh \mathcal{M}^k to the p -elements of a refined mesh \mathcal{M}^{k+1} through a *subdivision matrix* $\mathbf{R}^{[p],k}$, converging to a p -form in the limit of refinement. The original Whitney forms on cubical grids correspond to a particular choice of subdivision matrix $\mathbf{R}^{[p],k}$ between two consecutive mesh levels \mathcal{M}^k and \mathcal{M}^{k+1} [88], as illustrated in 2D in Fig. 3(middle). Since a given subdivision scheme can be turned into a subdivision scheme of higher regularity through convolution along grid directions [75], subdivision-based high-order Whitney forms can be derived from the original low-order Whitney forms, as long as the associated p - and $(p+1)$ -form subdivision schemes are treated in pairs to enforce that their subdivision matrices commute with the discrete exterior derivatives—thus ensuring a proper de Rham complex for the resulting finite element spaces spanned by these higher-order Whitney forms. For instance, in 2D, one such regularity-raising convolution results in bicubic splines for 0-forms and biquadratic splines for 2-forms as described in [84] (see Fig. 3(right)); the more convolutions one uses, the smoother the associated Whitney forms are and the larger their supports become. The higher-order Whitney p -form $\phi_i^{[p],k}$ associated with the i -th p -element σ_i^k at level k can thus be conceived as the limit of repeated subdivisions of an initial assignment of scalar values δ_{ij} on the p -elements σ_j^k , i.e.,

$$\phi_i^{[p],k} = \lim_{m \rightarrow \infty} \mathbf{R}^{[p],m} \mathbf{R}^{[p],m-1} \dots \mathbf{R}^{[p],k+1} \mathbf{R}^{[p],k} \mathbf{e}_i^k,$$

where the vector \mathbf{e}_i^k of real values on elements of the mesh \mathcal{M}^k contains a single 1 corresponding to element σ_i^k and zero on all other p -elements.

4.2.4. Properties of subdivision-based high-order Whitney forms

The resulting higher-order Whitney forms have also a few key distinguishing features that we summarize here, as these properties coming from their subdivision nature will be heavily exploited later on.

Simple degrees of freedom. While finite element exterior calculus methods use local polynomial shape functions associated with various degrees of freedom per mesh element [57, 80], our framework uses only the p -dimensional cells of a mesh to discretize a p -form. Consequently, the subdivision-based de Rham complex is simply the complex of the original low-order Whitney forms, i.e., the dual of the usual chain complex (see Fig. 2).

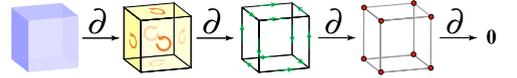


Fig. 2: **Chain complex.** Diagram depicting the chain complex of a cube element with the boundary operator: from the cube, to its square oriented faces, to their oriented edges, and to their vertices.

Finite-element spaces of differential forms. High-order Whitney forms $\phi_i^{[p],k}$ can still be used to construct differential p -forms from mesh values through Eq. (21); the associated finite element spaces of differential forms on the level- k mesh \mathcal{M}^k are, for $p=0, \dots, 3$,

$$\Lambda_p^k = \text{span} \left\{ \phi_i^{[p],k} \right\}_{i=1}^{n_k} = \left\{ \sum_{i=1}^{i=n_k} u_i \phi_i^{[p],k}(x) \mid u_i \in \mathbb{R} \text{ is associated to } p\text{-element } \sigma_i^k \in \mathcal{M}^k \forall i \right\}; \quad (24)$$

that is, they define “basis functions” for p -forms that use only degrees of freedom associated with p -dimensional elements, but are of higher regularity than Whitney’s original p -form basis, and with a larger spatial support. They can be used as test functions in our operator-adapted wavelet construction, i.e., we can set $\mathcal{V}_p^k := \Lambda_p^k$ for p -forms.

Refinement property. The built-in refinability of high-order Whitney p -forms implies

$$\phi_i^{[p],k} = \sum_{j=1}^{n_{k+1}} \mathbf{R}_{ji}^{[p],k} \phi_j^{[p],k+1}, \quad (25)$$

that is, if all the Whitney forms of level k are stored in a row vector $\Phi^{[p],k}$, one has $\Phi^{[p],k} = \Phi^{[p],k+1} \mathbf{R}^{[p],k}$.

Commuting with exterior derivative. If we denote by ∂_p^k ($p=0, \dots, 2$), the matrix of signed incidence between p - and $(p+1)$ -elements of mesh \mathcal{M}^k implementing the continuous boundary operator ∂ described in Fig. 2, then its transpose $\mathbf{D}_p^k := (\partial_p^k)^T$ represents the discrete exterior derivatives acting on finite-element forms and implements the continuous derivative \mathbf{d} [43], satisfying $\mathbf{D}_{p+1}^k \mathbf{D}_p^k = 0$ by construction. The subdivision matrices of high-order Whitney forms are chosen so that they commute with the discrete exterior derivatives [88]:

$$\mathbf{D}_p^{k+1} \mathbf{R}^{[p],k} = \mathbf{R}^{[p+1],k} \mathbf{D}_p^k.$$

That is, subdividing a discrete p -form on level k followed by application of the exterior derivative on level $k+1$ is the same as first applying the exterior derivative on level k followed by the subdivision scheme for $(p+1)$ -forms.

Fast pointwise and integral evaluations. Exact pointwise evaluation of the resulting high-order basis functions (which are piecewise polynomial for our regular meshes) at arbitrary parameter locations can be done efficiently [89]. Moreover, for each subdivision scheme based on a subdivision matrix $\mathbf{R}^{[p],k}$, there is an associated *limit stencil* operator $\mathbf{S}^{[p],k}$ defined through its matrix representation as:

$$\mathbf{S}_{ij}^{[p],k} = \int_{\sigma_i^k} \phi_j^{[p],k}. \quad (26)$$

With this limit stencil, one can evaluate the integral of a finite-dimensional p -form on all the p -elements of a mesh: indeed, from a vector of real values U^k on the p -elements $\{\sigma_i^k\}_{i=1}^{n^k}$ of mesh \mathcal{M}^k , the p -form $\mathbf{u} \in \Lambda_p^k(\mathcal{M})$ defined through

$$\mathbf{u}(x) = \Phi^{[p],k} U^k$$

integrated on a mesh element σ_ℓ^k is a local linear combination of the entries of U^k equal to the ℓ -th entry of the vector $\mathbf{S}^{[p],k} U^k$, that is,

$$\int_{\sigma_\ell^k} \mathbf{u} = [\mathbf{S}^{[p],k} U^k]_\ell.$$

Since the high-order Whitney forms are localized, $\mathbf{S}^{[p],k}$ is sparse; moreover, the non-zero values of this matrix can be evaluated in closed form through eigenanalysis of the subdivision matrix $\mathbf{R}^{[p],k}$ [75, 84], making the evaluation of integrals over mesh elements of any finite-element differential forms particularly efficient. Note that the original Whitney forms satisfy $\mathbf{S}_{ij}^{[p],k} = \delta_{ij}$ (see Eq. (22)) due to their histopolation property; hence, the values of the vector U^k are directly the integrals of the continuous form. This histopolation property is no longer true for a higher-order Whitney form, but its limit stencil encodes precisely the map between degrees of freedom and local integrals of the corresponding finite-element p -form.

Discrete Stokes' theorem. Due to the commutativity with discrete exterior derivatives and the existence of limit stencils, high-order Whitney forms induce a discrete variant of Stokes' theorem, $\int_{\mathcal{R}} \mathbf{d}\mathbf{u} = \int_{\partial\mathcal{R}} \mathbf{u}$ for any region \mathcal{R} and p -form \mathbf{u} : for a vector of real values U^k on the p -elements $\{\sigma_i^k\}_i$ of mesh \mathcal{M}^k ,

$$\mathbf{D}_p^k \mathbf{S}^{[p],k} U^k = \mathbf{S}^{[p+1],k} \mathbf{D}_p^k U^k. \quad (27)$$

In other words, Stokes' theorem holds over every mesh element in our finite-dimensional setting.

In the remainder of our exposition, we will often omit the order $[p]$ of the Whitney forms or even the level k over which it is defined to improve legibility, when this does not lead to any ambiguity. Therefore, high-order Whitney forms will often be simply referred to as ϕ .

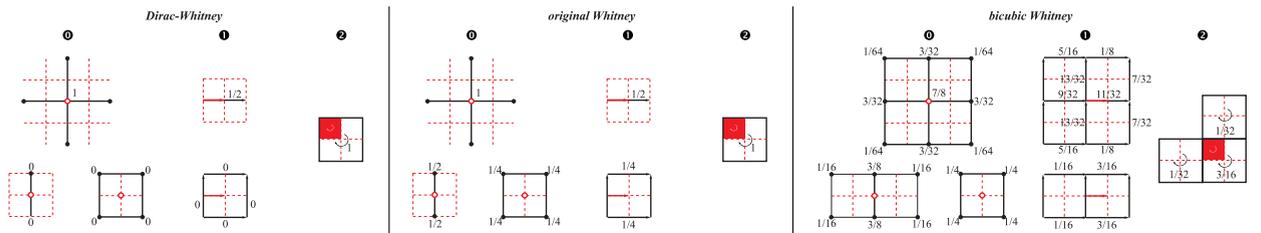


Fig. 3: **Refinement matrices for subdivision-based Whitney forms.** We provide the refinement rules in 2D and for 0-, 1- and 2-forms for (left) Dirac-Whitney, (middle) original Whitney, and (right) “bicubic” Whitney forms. For a k -element on a twice-refined grid, the refinement rules use a simple linear combination of nearby coarse elements to define a refinable k -form basis.

4.2.5. Finest discretization through Subdivision-based Exterior Calculus

Consider a continuous, linear, and bijective *exterior* operator \mathcal{L} that acts on differential p -forms and returns p -forms as well. In addition, we assume it is local, self-adjoint and positive-definite with respect to L^2 inner product on differential forms (formally defined as $\langle \mathbf{u}, \mathbf{v} \rangle_{L^2} = \int_{\Omega} \mathbf{u} \wedge \star \mathbf{v}$). Note that the bijectivity condition is needed to ensure the inverse is well-defined: in practice a differential operator is combined with suitable boundary conditions to give rise to a well-defined problem. Given such an operator, the differential equation $\mathcal{L}u = g$ must first be discretized on the finest mesh \mathcal{M}^q . Many finite element based discretizations can be applied for this step. The computational foundations of Finite Element Exterior Calculus [57], Discrete Exterior Calculus [43], and Mimetic (or Compatible) Operators [90] differ in how they deal with the discretization of PDEs: for instance, some define a discrete operator corresponding to

the Hodge star \star , while others prefer to bypass the creation of a discrete Hodge star by directly computing a discrete version of the codifferential $\delta = (-1)^{n(k-1)+1} \star \mathbf{d} \star$ to enforce adjointness with the exterior derivative \mathbf{d} ; boundary treatments are also often different as a consequence.

In our context, we remain agnostic vis-a-vis the discretization process of the weak form of the equation $\mathcal{L}u = g$. The only requirement is that our subdivision-based Whitney forms must be used as finite-element basis functions of differential forms, so that we can leverage their refinability with simple element-based degrees of freedom and structural properties (de Rham complex, Hodge decomposition, Stokes' theorem, etc.) to construct operator-adapted wavelets. For illustration purposes and numerical tests, we adopt the Subdivision-based Exterior Calculus framework (SEC [84]) based on the finite element spaces of differential forms $\{\Lambda_p^k\}_p$ defined in Eq. (24) to provide a discretization of exterior operators in this paper, but variants can be easily substituted.

Once this discretization is done on the finest level q , the rest of the operator-adapted multiresolution construction is mostly unchanged as we review next.

4.3. Operator-adapted wavelets

While the subdivision-based Whitney forms naturally lead to a L^2 -orthogonal notion of wavelets for any differential p -form by exploiting the difference between two consecutive approximation spaces $\mathcal{V}^k := \Lambda_p^k$ and $\mathcal{V}^{k+1} := \Lambda_p^{k+1}$ on two mesh levels \mathcal{M}^k and \mathcal{M}^{k+1} , the induced multiresolution decomposition leads to the same shortcomings as for the wavelet-Galerkin approach—but this time in the context of finite-element differential forms. In this section, we discuss how to construct operator-adapted wavelets in order to block-diagonalize a self-adjoint, positive-definite operator \mathcal{L} acting on differential p -forms.

Bottom-up construction of wavelet p -forms. Since the construction provided in Sec. 3 made no assumption on the choice of mesh hierarchy or basis functions, it applies nearly *as is* for arbitrary p -forms. Because the notion of “refinement” and “subdivision” are in fact equivalent (inasmuch as they both define linear relationships between functions at two consecutive levels), one simply has to:

1. set n_q to be the number of p -elements in the finest mesh \mathcal{M}^q ;
2. bootstrap the construction by replacing the finest scale basis functions φ_i^q by Whitney p -forms ϕ_i^q (of any chosen smoothness order) on mesh \mathcal{M}^q , and compute the corresponding stiffness matrix $\mathbb{A}^q := \mathbf{A}^q$;
3. use refinement matrices \mathbf{C}^k as the transpose of the subdivision matrices \mathbf{R}^k associated with the selected Whitney p -forms, i.e., $\mathbf{C}^k := \mathbf{R}^{k,T}$, which in turn impose wavelet refinement matrices \mathbf{W}^k .

Everything else remains *unchanged* if the indices used in the procedure are understood to refer to an enumeration of the p -cells of each mesh; in particular, the number N_k of wavelet coefficients at level k remains equal to $n_{k+1} - n_k$ as it represents the dimension of the finite-element p -forms that are in \mathcal{V}^{k+1} but not in \mathcal{V}^k . The resulting operator-adapted basis functions and wavelets will thus be linear combinations of Whitney p -forms on the finest grid, and can thus be stored as a sparse row vector per level as described in Sec. 3.6.

Choice of refinement matrix. As we just noted, a natural choice for a refinement matrix \mathbf{C}^k to be used in our operator-adapted wavelet construction is (the transpose of) the subdivision matrix \mathbf{R}^k associated to Whitney p -forms of a chosen regularity order. This includes the subdivision matrices of the original Whitney p -forms (see Fig. 3(middle)), as well as any of the higher order Whitney subdivision matrices (see an example in Fig. 3(right)). Since the refinement matrix describes how a form basis element at level k can be written as a linear combination of the form basis elements at level $k+1$, the sparsity of $\mathbf{C}^k = \mathbf{R}^{k,T}$ decreases with the smoothness order of the corresponding Whitney forms, in turn impacting the sparsity and locality of \mathbf{W}^k .

Remark 3: Observe that the multiresolution construction could use a given subdivision-based Whitney p -form at the finest scale, whose regularity impacts the quality of the discretization \mathbb{A}^q of operator \mathcal{L} , paired with an *arbitrary* refinement matrix \mathbf{C}^k . This choice of refinement, instead, impacts the regularity of test functions on coarser levels depending on the spatial size of the refinement stencil. As demonstrated in [32], approximation qualities of test functions in the weak H^{-s} -norm carry over to approximation qualities of the resulting operator-adapted wavelets in stronger energy norm. This means that using a simpler (i.e., sparser) refinement matrix \mathbf{C}^k often turns out to be sufficient to obtain a high-quality operator adaptation, while offering a simplified computational procedure, faster decaying basis functions, and better condition numbers as we will demonstrate later. On the other hand, such simplified refinement rule obviously weakens the homogenization effect and thus, reduces the efficiency of model reduction. Considering \mathbf{C}^k and \mathbf{R}^k decorrelated just offers added flexibility to our construction.

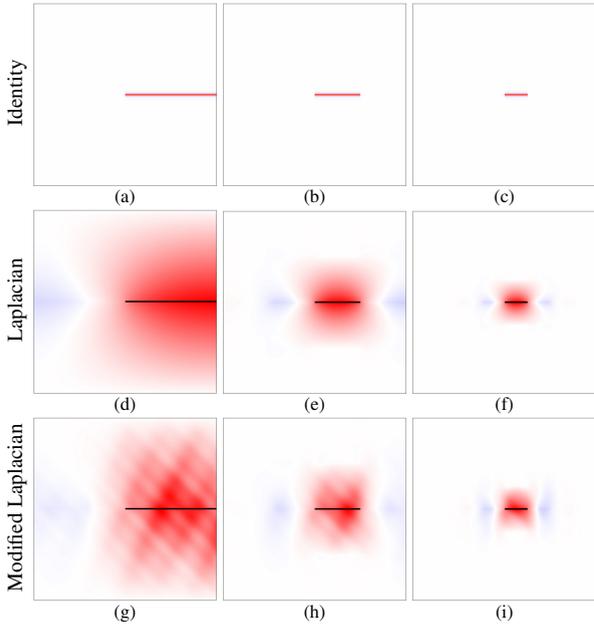


Fig. 4: **Dirac-Whitney basis functions adapted to elliptic operators.** Vertical components of 1-form basis functions φ_i^k adapted to different elliptic operators using the Dirac-Whitney refinement 1-form rule and associated with horizontal edges (in thick black) are visualized at three different resolution levels (2×2 , 4×4 , and 8×8 grids) using a linear red-to-blue colormap: (a)-(c) Identity operator; (d)-(f) 1-form Laplacian; (g)-(i) modified 1-form Laplacian (see Eq. (28)).

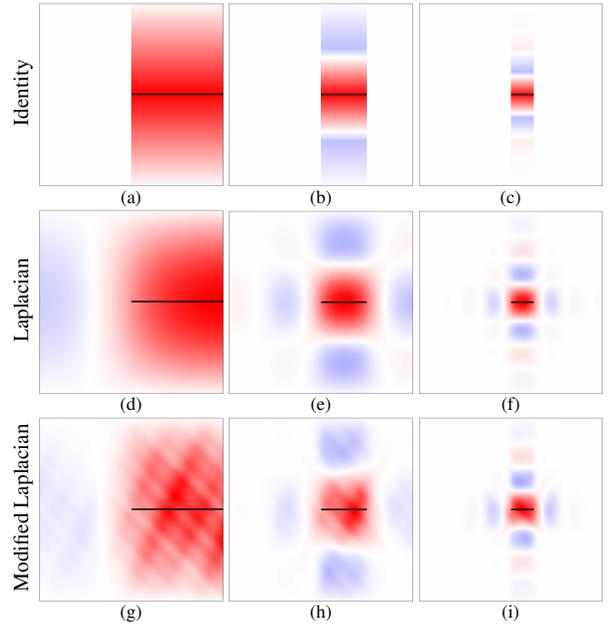


Fig. 5: **Whitney basis functions adapted to elliptic operators.** Vertical components of 1-form basis functions φ_i^k adapted to different elliptic operators using the original Whitney refinement 1-form rule and associated with horizontal edges (in thick black) are visualized at three different resolution levels (2×2 , 4×4 , and 8×8 grids) using a linear red-to-blue colormap: (a)-(c) Identity operator; (d)-(f) 1-form Laplacian; (g)-(i) modified 1-form Laplacian (see Eq. (28)).

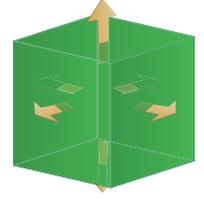
Dirac-Whitney refinement matrix. While the original (low-order) Whitney forms lead to the sparsest refinement matrices in the subdivision-based Whitney family, an even sparser choice of refinement matrix is worth mentioning: because our mesh hierarchy is nested, the p -elements of mesh \mathcal{M}^k can be written as linear combinations of p -elements of mesh \mathcal{M}^{k+1} . This linear relationship can be seen as the refinement rule for the basis of *chains* or, equivalently, of *cochains*. One can thus see these refinement matrices as defining a dual basis to currents: they are measure-based Whitney forms that include, e.g., a Dirac delta function per vertex for 0-forms, and an indicator function of each 3-cell for 3-forms. For this reason, we refer to these sparsest refinement matrices as *Dirac-Whitney refinement matrices* (see Fig. 3(left)). While the corresponding (generalized) test functions are not quite useful in numerical applications (and would clearly be a poor choice of test functions), our construction can, in fact, use these refinement matrices: when paired with any subdivision-based Whitney forms on the finest level q to bootstrap the construction as discussed in Rmk. 3 above, our resulting operator-adapted basis functions and wavelets will have very localized refinement matrices, but they will all be linear combinations of Whitney forms defined on mesh \mathcal{M}^q . As we will demonstrate in Sec. 5, this approach still leads to operator adaptation, with increased efficiency due to the sparsity of the Dirac-Whitney refinement matrices (see Fig. 4 vs. Fig. 5).

4.4. Divergence-free operator-adapted wavelets

Oftentimes, differential equations of the form of Eq. (1) must be solved *within a given linear subspace*. This is the case, for instance, for incompressible fluids or solids where the additional incompressibility condition $\nabla \cdot u = 0$ plays a crucial role in the resulting dynamics. Enforcing this constraint has a long history in finite elements [91]. We show here that our operator-adapted wavelet construction can accommodate this case quite simply as well. We first focus on the case in which original Whitney forms are used, before extending it straightforwardly to higher-order Whitney forms.

Discrete divergence. Stokes, Euler or Navier-Stokes equations in fluid dynamics are conveniently expressed with differential forms [58, 63, 60]. A typical discretization is to consider the velocity field as a 2-form in 3D, i.e., as a flux through each face of a 3D grid so that no-flux boundary conditions are trivially enforced [92]. A set of discrete

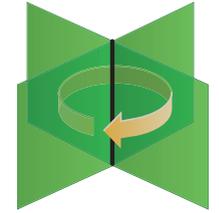
fluxes $U^q = (u_1^q \dots u_{n_q}^q)^T$ on the finest resolution level \mathcal{M}^q are converted into a continuous form \mathbf{u} using the face-based low-order Whitney basis functions $\{\phi_i^q\}_{i=1}^{n_q}$, where n_q is the total number of 2-cells in \mathcal{M}^q . Note that the integral of the resulting continuous 2-form \mathbf{u} on a face is precisely the flux initially stored on that face due to the histopolation property of linear Whitney forms (see Eq. (22)). The 2-form \mathbf{u} (or, equivalently, the corresponding vector field $(\star\mathbf{u})^\sharp$ obtained by lowering the indices of the form once a metric is given) is divergence-free if and only if $\mathbf{d}\mathbf{u} = \mathbf{0}$, where \mathbf{d} is the exterior derivative operator. In our finite dimensional setting, it simply implies that the sum of the (oriented) fluxes leaving any 3-cell of the mesh (see inset) is zero, which is denoted as $\mathbf{D}_2^q U^q = \mathbf{0}$, where \mathbf{D}_2^q is the discrete exterior derivative operator (the transpose of the boundary (incidence) operator acting on 3-cells, see [43] for instance) at level q .



Divergence-free test forms through stream functions. Instead of enforcing the linear divergence-free constraint on U^q , one can instead work in the reduced space of divergence-free vector fields. Because the whole setup of finite-element exterior calculus offers a discrete de Rham complex and since we assume our domain to be of trivial topology (no harmonic terms), we can consider instead a set of 2-forms, one per edge of \mathcal{M}^q , that spans this reduced space. This basis of divergence-free 2-form is defined as

$$\Xi^q = \Phi^q \mathbf{D}_1^q,$$

where Ξ^q denotes the row vector of edge-based 2-forms $\{\xi_i^q\}$ and Φ^q is, similarly, the row vector of all face-based Whitney 2-forms $\{\phi_i^q\}$. Hence, each 2-form ξ_i^q is a linear combination of local Whitney 2-forms with coefficients based on the discrete exterior derivative \mathbf{D}_1^q , a matrix whose rows and columns correspond to faces and edges of mesh \mathcal{M}^q respectively. For instance, for a regular 3D grid, each 2-form ξ_i^q associated with the i -th edge is simply the signed sum of the Whitney basis functions on the four faces adjacent to the edge (see inset). By construction, each edge-associated 2-form is divergence-free since \mathbf{D}_1^q is in the kernel of \mathbf{D}_2^q (i.e., $\mathbf{D}_2^q \mathbf{D}_1^q = \mathbf{0}$ corresponding to the continuous property $\mathbf{d}\mathbf{d} = \mathbf{0}$). Hence, the finite dimensional vector space spanned by $\{\xi_i^q\}_i$ (they form a basis of this space due to their linear independence) represents the set of all 2-forms (or equivalently, vector fields) that are divergence-free in the discrete sense. This is nothing else but the differential form equivalent of defining the velocity through the curl of a *stream (vector-valued) function*: the degrees of freedom are thus edge-based values $\{s_i\}_i$ representing the local integrals of the stream function along edges, while $\sum_i s_i \xi_i$ is the discrete divergence-free 2-form.



Operator-adapted divergence-free wavelets. The bottom-up approach described in Sec. 4.3 can now be applied directly with the divergence-free 2-form basis Ξ^q used as test functions at the finest scale to bootstrap the construction. The only difference is that n_q is now the number of *edges* (since it represents the number of degrees of freedom in this divergence-free case) and one must use a refinement matrix $\mathbf{C}^{(1),k}$ corresponding to 1-forms, even though we are constructing 2-form wavelets. With this approach, we work directly in the space of divergence-free fields without the need for reprojection steps: both operator-adapted bases and wavelets will be divergence-free by construction, since they are linear combinations of the elements of $\{\xi_i^q\}_{i=1}^{n_q}$. Note that in 2D, the stream function is a 0-form and fluxes through edges are 1-forms, so the construction must be altered to use linear combinations of edge-based basis functions per vertex instead, but the overall approach is identical otherwise.

Higher order bases. While the divergence-free construction we just described assumes that the refinement matrices corresponding to the original Whitney forms are used, the same construction applies to higher order Whitney forms as well. Indeed, when using subdivision-based Whitney bases, the integral of the reconstructed continuous 2-form \mathbf{u} over an oriented face is always a local linear combination of the initial face values $\{u_i^q\}_i$, called the limit stencil $\mathbf{S}^{[2],q}$ (see Eq. (26)). Therefore, the vector of divergences over all 3-cells is given by $\mathbf{D}_2^q \mathbf{S}^{[2],q} U^q$, which by Eq. (27) (discrete Stokes' theorem), is also equal to $\mathbf{S}^{[3],q} \mathbf{D}_2^q U^q$. Enforcing $\mathbf{D}_2^q U^q = \mathbf{0}$ thus implies discrete divergence-freeness of the resulting 2-form, and the construction carries over as is, even in this high-order case.

Non-trivial topology. If the domain has non-zero genus and requires a non-regular grid, the space of divergence-free forms/fields needs to account for harmonic forms. Thankfully, computing these additional basis functions is a well studied problem which can be achieved by leveraging the homology generators of the domain, see for instance [43]. As long as all the meshes \mathcal{M}^k in the hierarchy have the same number of boundaries and the same genus (that is, they all capture the proper topology of the domain), these extra topological degrees of freedom are present at each level and are thus trivially added to the reduced bases of divergence-free forms at each level.

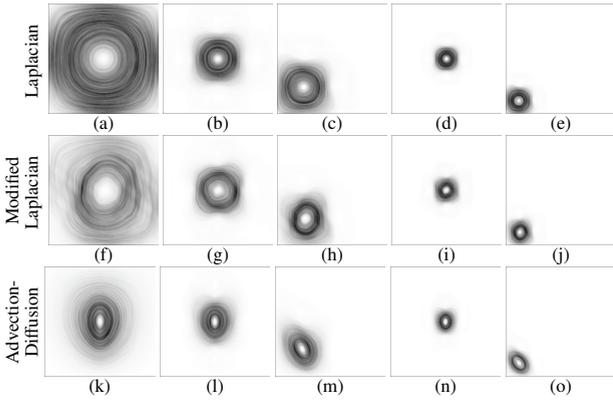


Fig. 6: **Operator-adapted divergence-free bases.** We visualize vector fields corresponding to divergence-free 1-form basis functions φ_i^k adapted to various operators and associated to several spatial locations on three different resolution levels. Divergence-free basis functions retain a characteristic vortical appearance; however, the differences in their shapes reflect their ability to capture representative features of the associated solution spaces: they are smooth for the 1-form Laplacian (top), have high frequency oscillations for the modified 1-form Laplacian (middle, see Eq. (28)), and stretched along the advecting velocity field (from Fig. 19(left)) for the advection-diffusion operator (bottom).

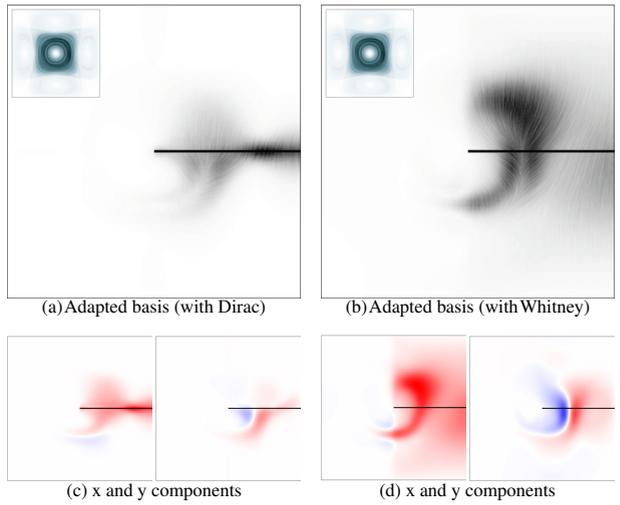
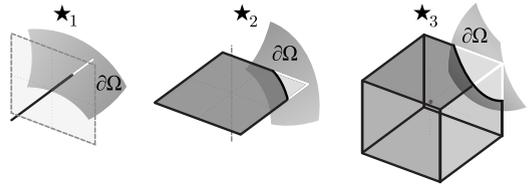


Fig. 7: **Adapted basis for advection-diffusion.** We visualize a coarse-edge (in thick black) basis function φ_i^k adapted to advection-diffusion, where the advecting velocity field is shown in the top-left inset. Instead of the first-order upwind discretization used in other figures, an energy-preserving discretization of the advection operator [93] is employed here.

4.5. Embedding complex domains on Cartesian grids

While unstructured meshes can conform to arbitrary domains, regular grids require simpler data structures and simpler refinement stencils which make them highly desirable when efficiency is paramount. However, handling complex domains clashes with the simplicity of regular grids.

Recently, in the context of DEC discretization of differential calculus, a local numerical homogenization of the diagonal Hodge star to capture sub-grid resolution was proposed in [61]. It extended to arbitrary k -forms the approach of [94], which offered a robust and second-order convergent pressure projection over an arbitrary domain Ω using a regular computational grid. While standard diagonal entries of the standard diagonal Hodge p -star are defined as the ratios of volumes of primal p -element and its dual element [54], the key idea of the improved boundary treatment is to account for the parts of the regular grid that are outside the domain Ω by altering the discrete diagonal Hodge star accordingly: partial lengths, areas, or volumes of primal elements that are within the domain Ω are used in the evaluation of the diagonal Hodge star, see inset. Topological operators like the exterior derivative remain unchanged. This method can be used in our context as well: one can construct arbitrary operator-adapted basis functions and wavelets on a regular grid while conforming to a smaller domain Ω by modifying the Hodge star operator on the finest mesh \mathcal{M}^q : as Figs. 8 and 9 demonstrate, an operator \mathcal{L} involving Hodge stars in its exterior calculus expression will be properly handled over the actual domain Ω as long as it is covered by the computational grid. Note that this approach, which *homogenizes both space and operator* over the regular grid, is currently limited to a low-order approximations of the Hodge star. Future extensions to higher-order approximations of Hodge stars may provide more accurate sub-grid accuracy.



4.6. Discussion

A few remarks are in order about our extension to multiresolution analysis of differential forms.

L²-adapted case. When the identity operator is used along with Whitney p -forms at the finest level and their corresponding refinement matrices, one may expect that the resulting operator-adapted basis functions φ^k will simply be the same Whitney p -forms on each p -element of the mesh \mathcal{M}^k at each level k . However, as Rmk. 2 in Sec. 3.1 pointed out, this would only be true if these Whitney forms were L^2 -orthonormal in the first place, which is not the case. However, the induced L^2 -orthonormal basis functions φ^k share a close resemblance to the Whitney forms from which the refinement matrix derives, as Fig. 5 demonstrates.

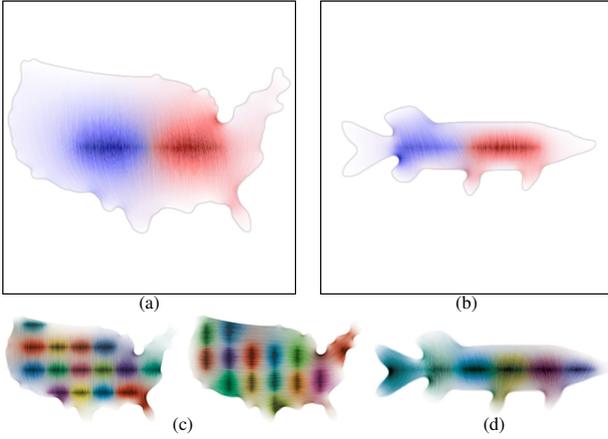


Fig. 8: **Laplace-adapted basis functions on grid-embedded domains.** Discretization of the 1-form Laplace operator on a regular grid can be adjusted to account for an arbitrary computational domain, like a US map or a fish: using a finest grid of 256×256 elements, two resulting coarse (4×4 level) operator-adapted basis functions φ_i^k (one in red, one in blue) are displayed in (a-b), and all the basis functions at the level for which the grid is 8×8 are shown using different colors in (c-d).

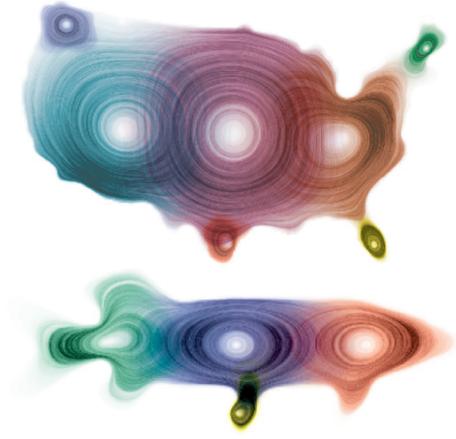


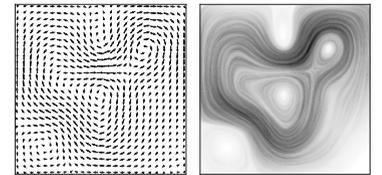
Fig. 9: **Divergence-free adapted basis functions on grid-embedded domains.** Using the same setup as Fig. 8, we now show the divergence-free basis function adapted to the 1-form Laplacian restricted to complex embedded domains. Various basis functions on a coarse (4×4) and a finer (8×8) level are shown using different colors in: three coarse and four fine basis functions for the US map, and three coarse and one fine basis functions for the fish.

Impact of refinement matrix choice. Having the finest basis functions used to bootstrap the construction and the actual refinement matrix \mathbf{C}^k being decorrelated as noted in Rmk. 3 of Sec. 4.3 offers much flexibility in practical computations. Dirac-Whitney refinement matrices will have operator-adapted basis functions and wavelets with the smallest support, and this support will grow with the smoothness order of the Whitney forms that the refinement matrix corresponds to. Higher-order Whitney forms are advisable from a functional approximation point of view as they induce lower approximation errors for a given resolution level. If, instead, lowering computational cost is paramount, then low-order Whitney forms are preferable. In particular, Dirac-Whitney refinement matrices typically induce lower condition numbers of matrices \mathbf{A}^k and \mathbf{B}^k : their associated refinement kernel matrix \mathbf{W}^k have optimal condition numbers, i.e., $\text{cond } \mathbf{W}^k = 1$, while the use of higher-order Whitney forms leads to higher (but still bounded) condition numbers.

Linear differential constraints. While we only gave details for the case of divergence-free vector fields, the construction of a basis of differential forms in a given linear subspace is possible for many typical differential constraints in physical systems: one only has to find a set of localized linear combinations of Whitney forms that creates a basis of the kernel of the linear operator at play in the constraint.

5. Applications and Numerical Tests

In this section, we discuss potential applications of our operator-adapted decomposition of differential forms and present a number of numerical experiments to illustrate our contributions. Since our work extends the recent work of Owhadi *et al.* [28, 29, 30, 31, 32] on operator-adapted decompositions of scalar functions, we focus on examples involving one-forms (instead of zero-forms) to highlight the novelty of our work, even if our approach applies to arbitrary forms. We will depict one-forms using a variant of Line Integral Convolution (LIC [95], implemented in [96]), which offers a more expressive visualization than the traditional quiver plot of vectors (see inset) through the use of integral lines and greylevel coloring to encode direction and length respectively.



Quiver plot vs. LIC visualization

5.1. Elliptic operators

We first show examples of operator-adapted basis functions and wavelets for typical linear operators. In these three examples, the original Whitney 1-forms are used at the finest level to bootstrap the decomposition.

Identity operator. As a trivial example, we can run our algorithm on 1-forms for $\mathcal{L} = \mathbb{I}$. As expected, using a Dirac-Whitney refinement matrix leads to nearly edge-based Dirac basis functions, while using a low-order Whitney refinement matrix engenders an L^2 -orthogonal basis of 1-forms resembling the edge-based Whitney forms, see Figs. 4 and 5 (top row).

Laplace–de Rham operator. For the 2D Laplacian $\mathcal{L} = \star_1 d_0 \star_0^{-1} d_0^T \star_1 + d_1^T \star_2 d_1$ of one-forms with zero Dirichlet boundary conditions (corresponding to the vector Laplacian $\nabla \nabla \cdot - \nabla \times \nabla \times$ in vector calculus), the adapted basis functions look much smoother (see Figs. 4 and 5 (middle row)). This is to be expected due to the obvious link between the Dirichlet energy (a common measure of smoothness) and the Laplace–de Rham operator: as indicated by the variational definition in Eq. (20), they are, in fact, the *smoothest* localized functions.

General elliptic operator. Our approach applies to any elliptic operator, such as the general diffusion operator Δ_σ , where the vector Laplacian (with zero Dirichlet boundary condition) uses an arbitrary conductivity tensor σ . We demonstrate this property by using $\mathcal{L} = \tilde{\star}_1 d_0 \star_0^{-1} d_0^T \tilde{\star}_1 + d_1^T \star_2 d_1$ on one-forms, where the Hodge star of 1-forms is altered by making it dependent on a spatially-varying metric μ of the same form as proposed in [50]:

$$\mu(x, y) = \begin{pmatrix} \alpha(x, y) & 0 \\ 0 & 1 \end{pmatrix} \quad \text{with} \quad \alpha(x, y) = \prod_{k=0}^4 \left(1 + \frac{1}{5} \cos(2^k \pi(x + y)) \right) \left(1 + \frac{1}{5} \sin(2^k \pi(x - 2y)) \right). \quad (28)$$

We denote the resulting modified Hodge star as $\tilde{\star}_1$. Figs. 4 and 5 (bottom rows) show the resulting operator-adapted 1-form basis functions, reflecting the highly-oscillatory modification from the vector Laplacian that the metric μ creates. We will show in Sec. 5.5 that these basis functions allow proper homogenization of this elliptic operator and lead to bounded stiffness condition numbers. One can thus see our decomposition as an extension of the numerical methods for scalar-valued homogenization of the operator $\nabla \cdot \sigma \nabla$, such as [50].

5.2. Advection-diffusion operator

Our approach can handle non-elliptic operators just as well: we demonstrate it by constructing a wavelet decomposition adapted to the well-studied advection-diffusion equation (with zero boundary flux and free-slip along the boundary):

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{a} \cdot \nabla \mathbf{u} - \nu \Delta \mathbf{u} = \mathbf{b}$$

where $\mathbf{a}, \mathbf{b} \in \mathbb{R}^2$ are given vector fields, $\nu \in \mathbb{R}_+$ is a viscosity parameter and $\mathbf{u} \in \mathbb{R}^2$ is the unknown, time-varying velocity field. In general, this equation describes the transport and diffusion of physical entities such as the ozone in the atmosphere or pollutants in oceans. Using its exterior calculus equivalent for one-forms (i.e., $u = \mathbf{u}^\flat$), a trapezoidal-based time discretization for a time step τ leads to the following implicit scheme (using $u_{n+\frac{1}{2}}$ as a shorthand for $\frac{1}{2}(u_{n+1} + u_n)$):

$$\frac{u_{n+1} - u_n}{\tau} + \mathfrak{L}_\mathbf{a} u_{n+\frac{1}{2}} - \nu \Delta u_{n+\frac{1}{2}} = \mathbf{b}^\flat \quad \implies \quad \left[\mathbb{I} + \frac{1}{2}\tau \mathfrak{L}_\mathbf{a} - \frac{1}{2}\tau\nu \Delta \right] u_{n+1} = \left[\mathbb{I} - \frac{1}{2}\tau \mathfrak{L}_\mathbf{a} + \frac{1}{2}\tau\nu \Delta \right] u_n + \mathbf{b}^\flat$$

where the operator $\mathfrak{L}_\mathbf{a}$ encodes the (Lie) advection of a form by the vector field \mathbf{a} . Observe that since the continuous form of the advection-diffusion equation mixes the skew-symmetric advection operator and the self-adjoint diffusion operator, the operator $\ell = \left[\mathbb{I} + \frac{1}{2}\tau \mathfrak{L}_\mathbf{a} - \frac{1}{2}\tau\nu \Delta \right]$ is not necessarily symmetric and positive-definite. We thus compute the basis functions adapted to the symmetrized operator $\mathcal{L} = \ell^T \ell$ instead. In our results, we discretize the advection-diffusion operator ℓ using a first-order upwind approach [97] for most numerical tests involving advection-diffusion; the only exception is in Fig. 7 where we use the antisymmetric and conservative discretization provided in [93], which leads to similar results once symmetrized: the adaptation to the operator \mathcal{L} manifests itself in the “adverted” shape of the resulting basis functions, see Figs. 6, 7 and 19.

5.3. Embedded domains

By exploiting the domain adaptation of the Hodge star explained in Sec. 4.5, we can construct operator-adapted basis functions over complex domains while still using Cartesian computational grids. Figs. 8 and 9 show vector-Laplacian-adapted edge functions over fish-shaped and US-shaped domains. Far away from the boundary, the edge functions are virtually unchanged, but they nicely adapt to the convoluted boundary for surrounding edges. While more precisely conforming edge functions could be constructed from a triangle or quadrangle mesh approximation of the shape, our sub-grid accurate treatment of arbitrarily-shaped domains over regular grids is trivial to implement and properly converges as the grid is refined [94, 61].

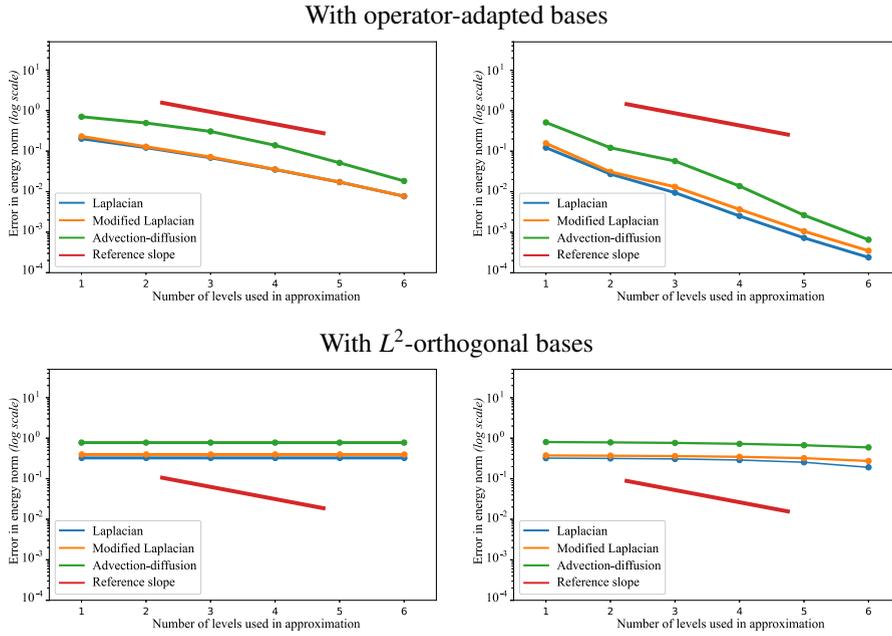


Fig. 10: **Homogenization Effect.** We demonstrate decay of the homogenization error — i.e., the energy norm of the approximation error between the fine FEM solution u and its level- k operator-adapted approximation u^k (see Eq. (A.1)) — as a function of k ; for each of the test operators (1-form Laplacian, modified 1-form Laplacian and advection-diffusion), the solution u is visualized in Fig. 1, to which the u^k approximant is compared. The red line depicts the optimal slope of the approximation error obtained by using eigenfunctions as basis functions (the so-called Kolmogorov n -width). Our operator-adapted Dirac-Whitney (top left) and original Whitney (top right) wavelets lead to exponential decay of the error for the three operators, demonstrating the expected numerical homogenization effect. In contrast, operator-independent wavelets based on Dirac-Whitney (bottom left) and original Whitney (bottom right) refinement rules fare significantly worse, both leading to unconvincing decay; these wavelets are L^2 -orthogonal between scales, i.e., adapted to the Identity operator, see Figs. 4 & 5 (top) for their associated basis functions.

5.4. Divergence-free basis functions

The construction of divergence-free vector-valued basis functions is an important computational need in many practical applications. While multiscale analysis approaches using divergence-free L^2 -orthogonal wavelets have been proposed (in both the tensor and non-tensor product cases) [98, 99], none have been offered in the more general case of operator-orthogonality. As detailed in Sec. 4.4, our approach lends itself quite straightforwardly to this case. We demonstrate the divergence-free basis functions adapted to three differential operators in Fig. 6 on a simple domain. We also leverage the domain adaptation of the Hodge star in Fig. 9 to provide divergence-free bases on complex domains: using only a Cartesian grid, we construct a hierarchy of divergence-free vector-Laplacian-adapted basis functions that are restricted to a spatial domain with complex boundaries and with zero normal components at the boundary. One can clearly see the effects of the domain in the way the basis functions associated to edges of various hierarchy levels are shaped through the domain.

5.5. Approximation properties

We demonstrate the coarse-graining properties of our operator-adapted basis functions through the analysis of approximation errors. The notion of numerical homogenization is characterized by the following bound on the energy norm of the difference between a solution u to Eq. (1) and its operator-adapted level k approximation u^k (see Table 1 for definition):

$$\|u - u^k\|_{\mathcal{L}} \leq 2^{-ks} C \quad (29)$$

for the case of our 2D regular mesh hierarchy, where the constant $C > 0$ does not depend on the choice of operator (see Appendix A or [28] for the general case).

Using high resolution FEM solution of Eq. (1) as a proxy for u , and the right hand side of Eq. (29) as a reference slope, we plot the decay of the energy norm of approximation error as a function of k in Fig. 10 (top row) for several examples. The slopes of resulting curves follow the theoretical bound of Eq. (29) for three test operators, with original Whitney refinement rule leading to a faster error decay compared to Dirac-Whitney. If non-adapted wavelets are used instead (see Fig. 10 (bottom row)), the approximation errors deteriorate significantly and even display some growth

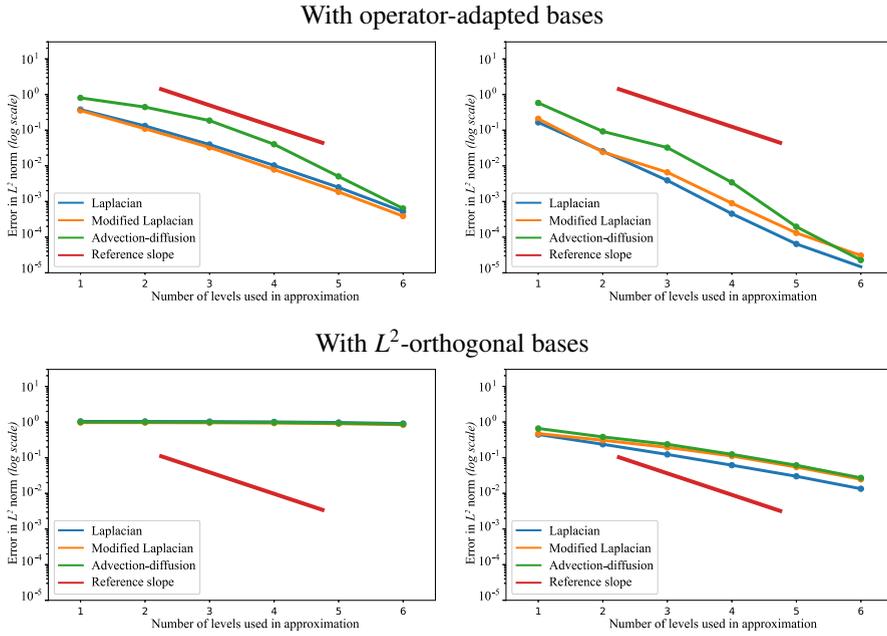


Fig. 11: **Approximation Errors.** Using the same setup as Fig. 10, we now measure the approximation error using the L^2 norm instead of the energy norm. With this norm, it is known that non adapted bases may also exhibit convergence, but their rate of convergence can be arbitrarily bad [100]. While the decay rate is improved for all curves compared to Fig. 10, operator-adapted wavelets (top row) still perform significantly better than non-adapted ones (bottom row); Dirac refinement rules (left column) lead to slower decay than Whitney (right column), as expected due to smaller stencil. Note that the red curve representing the theoretical rate of convergence here is two times steeper than before.

with k for Dirac-Whitney refinement rule. For completeness, we also provide a similar plot in Fig. 11, where now the L^2 -norm is employed to measure approximation error instead of the energy norm. Operator-adapted construction still exhibits much faster decay compared to non-adapted version.

The coarse-graining properties of our basis-functions imply that only a subset of adapted wavelets can be used to produce a high-quality approximation of the solution space of their associated operator, thus naturally opening possibilities for efficient model reduction. The visualization of multiresolution decomposition of u^q for the 1-form Laplace operator in Figs. 17 and 18 also confirms that insight and reiterates that finer resolution levels correspond to high-frequency components of the solution, albeit they look quite different for the original Whitney vs. the Dirac-Whitney refinement rules: the former has better homogenizing properties due to larger support, which results in its ability to carry over more information to coarser resolutions. Indeed, using only the 3 coarsest levels of Whitney wavelets leads to a smaller approximation error than using the 5 coarsest levels of the Dirac-Whitney construction; the Whitney-based construction allows to capture 99.9% of energy content with only 112 adapted wavelets, while Dirac-Whitney needs 1984 adapted wavelets to achieve the same energy content; for reference, a full solution u^q has 32512 degrees of freedom.

We finally note that our level k approximations u^k are in fact optimal both in Galerkin and game-theoretical sense, see Appendix B for more details.

5.6. Eigen ranges and condition numbers

An important measure of performance for our multiscale construction involves a spectral analysis of the resulting wavelet stiffness matrices \mathbb{B}^k . In Fig. 12, we show that the condition numbers of the stiffness matrices \mathbb{B}^k are much smaller than their counterparts \mathbb{A}^k —from two to six orders of magnitude for high frequencies. Moreover, in Fig. 13 we also show that the ranges of eigenvalues of the matrices \mathbb{B}^k overlap for consecutive values of k , reconstructing the full eigenvalue range of the input stiffness matrix \mathbb{A}^q (evaluated with non-adapted basis on the finest resolution q), confirming the proper capture of the operator solution space at various spatial resolutions. Note that performing the same numerical tests for divergence-free bases lead to very similar results, see Fig. 14.



Fig. 12: **Condition numbers.** We compare the condition numbers (in log scale) of stiffness matrices A^k of 1-form basis functions (in blue) and their associated wavelets $B^{k,l}$ (in orange) for our operator-adapted multiresolution decomposition with $q=7$ resolution levels (finest grid is 128×128 cells) for three operators: 1-form Laplacian (top row), modified 1-form Laplacian (middle row, see Eq. (28)), and advection-diffusion (right row). The Dirac-Whitney refinement rule (left column) tends to produce slightly better condition numbers for matrices $B^{k,l}$ compared to the original Whitney rule (right column) as expected, but in both cases they remain bounded and often peak on intermediate resolution levels. In particular, condition numbers of the wavelet stiffness matrices $B^{k,l}$ are up to 3 orders of magnitude smaller than those of A^k for large k (i.e., on finer scales).

5.7. Exponential decay of basis functions

As we discussed in Sec. 3.5, our operator-adapted basis functions decay exponentially fast as long as some mild assumptions on the interaction between test functions and operator hold (see Appendix A). Figs. 15 and 16 corroborate this statement in the case of exterior operators, even if divergence-free constraints are imposed.

5.8. Sparsity patterns

We finally provide an evaluation of the sparsity of the stiffness matrices resulting from our multiresolution construction. In Fig. 20, we visualize $\text{diag}[A^1, B^1, \dots, B^6]$, a principal submatrix of the global multiresolution matrix L (Eq. (5)) using the Laplace operator on 1-forms. A straightforward implementation of Alg. 2 using an off-the-shelf linear solver combined with a pruning of all entries of B^k that are less than $1e^{-15}$ in absolute value results in already sparse matrices. However, our fast solve through spatial localization produces even sparser matrices with a limited impact on accuracy: the relative L^2 error in the resulting operator-adapted basis functions ψ_i^k introduced by this sparsification does not exceed 0.75% (as discussed in Sec. 3.6, we used 3-element-wide neighborhoods localization for Dirac-Whitney refinements, and 4-element wide for Whitney to account for larger support). Observe that in this example using 1-forms, each B^k is represented by two blocks corresponding to horizontal and vertical edges respectively.

5.9. Examples of applications

Our construction of fine-to-coarse operator-adapted basis functions and wavelets have many potential applications in numerical contexts. We cite a few important examples next.

Model reduction. Operator-adapted basis functions are particularly relevant in numerical simulation: as we discussed early on, they offer a finite-dimensional basis of locally-supported functions that captures the eigenspace of the corresponding operator as well as possible. Therefore, they are particularly relevant in the context of model reduction and adaptive refinement, much more so than classical wavelets. We discuss next a few approaches to leverage our

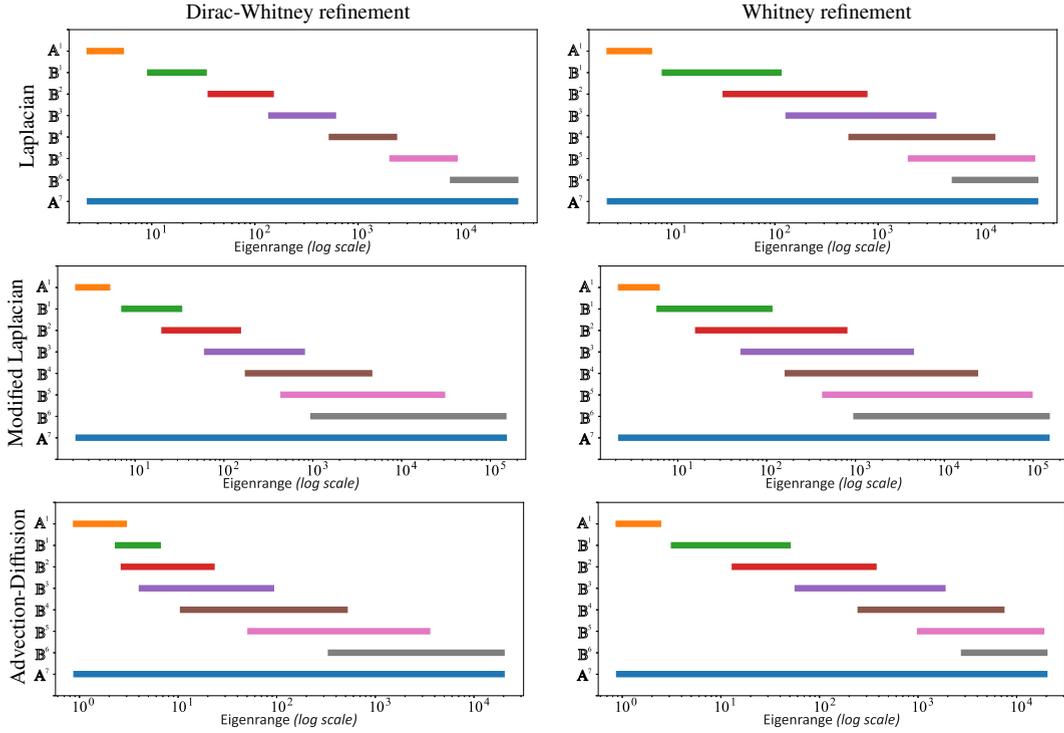


Fig. 13: **Eigenranges.** We show that the eigenvalue ranges of 1-form wavelet stiffness matrices \mathbb{B}^k are (overlapping) subbands of the eigenvalue range of the input stiffness matrix $A^l = A^l$ (evaluated with non-adapted basis on the finest resolution) for three test operators: 1-form Laplacian (top row), modified 1-form Laplacian (middle row, see Eq. (28)) and advection-diffusion (bottom row); each horizontal segment covers the spectrum of the associated stiffness matrix (A^l , B^1 , B^2 , B^3 , B^4 , B^5 , B^6 , and A^7). Both Dirac-Whitney (left column) and Whitney (right column) refinement rules produce operator-adapted wavelets that capture frequencies associated with respective resolution levels; the degree of overlap generally depends on the operator and tends to be smaller for Dirac-Whitney refinement due to its small spatial support.

operator-adapted multiresolution analysis to either improve accuracy for the same number of degrees of freedom, or reduce the number of degrees of freedom for a given accuracy.

- *Resolution reduction.* The simplest way to reduce the number of degrees of freedom in the discretization of a differential equation with limited impact on accuracy is to omit high frequencies (i.e., the “details” of the solution corresponding to finest operator-adapted wavelets) while retaining wavelets spanning the p coarsest resolution levels for $p < q$. Since our bottom-up construction accumulates information from finer scales to derive coarser basis functions, this simple approach allows to capture most of the solution without a large amount of fine scale wavelets.
- *Goal-driven model reduction.* While the previous approach is intuitive and simple, it fails at providing a clear assessment of the error induced by the removal of the finest details. Goal-oriented adaptive refinement techniques aim, instead, at selecting a subset of wavelets from each resolution level leading to a computationally efficient and high quality estimation of a particular output functional [101], typically defined as a linear functional Q . This is easily achieved in our framework: after computing wavelet coefficients w^k on level k via a sparse linear solve, the least contributing wavelets $\{\psi_j^k\}_{j=1}^{N_k}$ are found by checking if

$$|Q(w_j^k \psi_j^k)| \leq \varepsilon \max_{\substack{1 \leq p \leq k \\ 1 \leq i \leq N_p}} |Q(w_i^p \psi_i^p)|$$

for a small $\varepsilon > 0$. These wavelets can be safely eliminated to reduce the number degrees of freedom and make the matrices \mathbb{B}^k sparser while insuring that the quantity of interest Q can still be well approximated. Note that the error in Q incurred by removing the weakly contributing wavelets can be evaluated on the fly.

- *Norm-driven reduction.* Yet another strategy to drop degrees of freedom is to remove those that least contribute to the operator-norm of the solution. Since our wavelets are well localized, interactions between them tend to

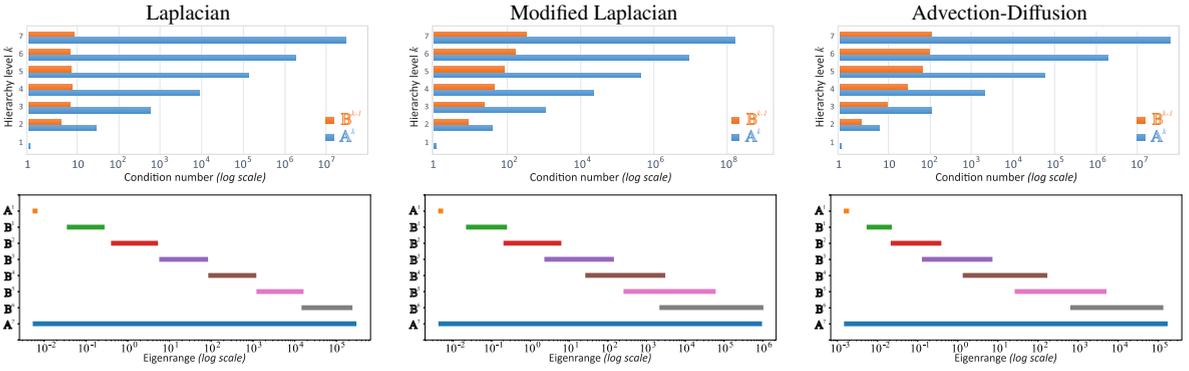


Fig. 14: **Eigenanalysis of operator-adapted divergence-free basis.** We demonstrate spectral properties of stiffness matrices corresponding to the 1-form divergence-free wavelets and basis functions adapted to 1-form Laplacian (left column), modified 1-form Laplacian (middle column, see Eq. (28)) and advection-diffusion (right column): condition numbers (top row) and eigenrange subband structure (bottom row) have the same qualitative properties as for unconstrained operator-adapted multiresolution constructions, see Figs. 12&13. Note that the range of eigenvalues for the top level is degenerate (hence the short orange bar): since there's only one vertex at the coarsest level, the stiffness matrix becomes just a scalar.

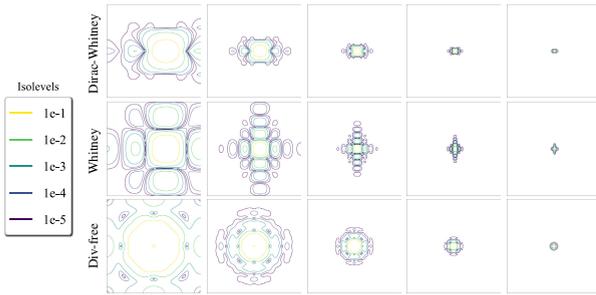


Fig. 15: **Localization of Laplacian-adapted bases.** We demonstrate the exponential decay of basis functions ϕ_i^k adapted to the 1-form Laplacian using Dirac-Whitney (top), original Whitney (middle) and divergence-free (bottom) refinements, normalized to have unit maximum value, using log-scale contour plots. Examples of edge-based basis functions from five different levels of the mesh hierarchy are shown.

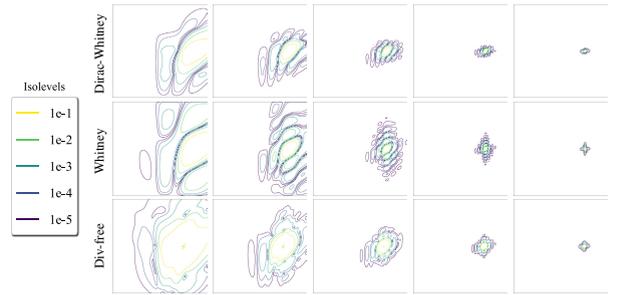


Fig. 16: **Localization of advection-diffusion-adapted bases.** We demonstrate the exponential decay of basis functions ϕ_i^k adapted to the advection-diffusion (discretized through first-order upwind approximation) using Dirac-Whitney (top), original Whitney (middle) and divergence-free (bottom) refinements, normalized to have unit maximum value, using log-scale contour plots. Examples of edge-based basis functions from five different levels of the mesh hierarchy are shown.

be limited. Therefore, a wavelet ψ_j^k at level k can be considered negligible in its impact on the solution if

$$\|w_j^k \psi_j^k\|_{\mathcal{L}}^2 \leq \varepsilon \max_{\substack{1 \leq \ell \leq k \\ 1 \leq i \leq N_i}} \|w_i^\ell \psi_i^\ell\|_{\mathcal{L}}^2$$

for a small $\varepsilon > 0$. However, since the operator norm of the solution

$$\|u^q\|_{\mathcal{L}}^2 = \sum_{i,j=1}^{n_1} u_i^1 u_j^1 \mathbb{A}_{ij}^1 + \sum_{k=1}^{q-1} \sum_{i,j=1}^{N_k} w_i^k w_j^k \mathbb{B}_{ij}^k$$

has non-negligible wavelet interaction terms associated with off-diagonal elements of the stiffness matrices, this simple heuristic can be made more robust by measuring the importance of a particular wavelet ψ_j^k through its full contribution to the norm, expressed as

$$\left| w_j^k \sum_{i=1}^{N_k} w_i^k \mathbb{B}_{ij}^k \right|.$$

Thresholding these contributions results in a well-approximated solution in the operator norm with typically very few coefficients.

Efficient Navier-Stokes simulation. Our approach also applies to nonlinear equations such as the Navier-Stokes equations for incompressible fluids:

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} - \nu \Delta \mathbf{u} &= 0 \\ \nabla \cdot \mathbf{u} &= 0 \end{aligned}$$

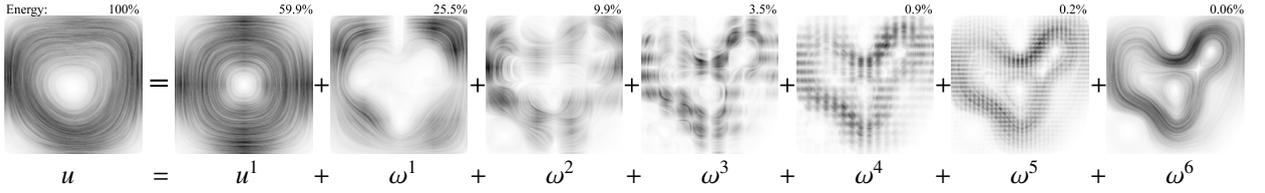


Fig. 17: **Operator-adapted multiresolution decomposition through Dirac-Whitney refinement.** For the 1-form Laplace operator, an element of its solution space 1-form u defined via edge values on a 128×128 grid (left, LIC-visualized as its equivalent vector field) can be efficiently decomposed into a sum of a coarse 1-form u^1 described via edge values on a 2×2 grid, and all the wavelet contributions $\omega^1, \omega^2, \omega^3, \omega^4, \omega^5, \omega^6$ of the mesh hierarchy. Since our vector visualization does not convey relative magnitude, we also indicate the relative energy content (as a percentage) of each component. Compare with Fig. 18.

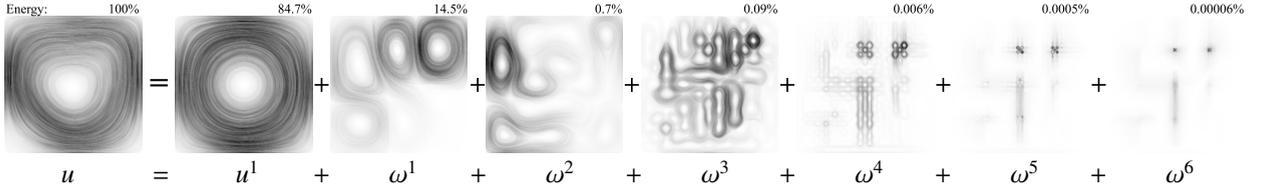


Fig. 18: **Operator-adapted multiresolution decomposition through original Whitney refinement.** For the 1-form Laplace operator, an element of its solution space 1-form u defined via edge values on a 128×128 grid (left, LIC-visualized as its equivalent vector field) can be efficiently decomposed into a sum of a coarse 1-form u^1 described on a 2×2 grid, and all the wavelet contributions $\omega^1, \omega^2, \omega^3, \omega^4, \omega^5, \omega^6$ of the mesh hierarchy. Since our vector visualization does not convey relative magnitude, we also indicate the relative energy content (as a percentage) of each component. Compare with Fig. 17.

Multiresolution analysis of incompressible fluids is an important topic in computational fluid dynamics, but an efficient construction of wavelets that are truly adapted to the fluid dynamics has remained elusive. Our construction offers such a tool, and could be extremely useful to turbulence analysis and efficient simulation.

For simple integrators where the update rule is linear in the next velocity \mathbf{u}^{n+1} , our approach can be used *as is* where the operator corresponds to the update rule just like we showed in Sec. 5.2 in the case of advection-diffusion. For instance, using a trapezoidal temporal discretization with operator splitting, we obtain the following Implicit-Explicit numerical scheme involving the usual pressure projection step:

$$\begin{aligned} \left[\mathbb{I} + \frac{1}{2}\tau \text{Adv}(\mathbf{u}_n) - \frac{1}{2}\tau\nu\Delta \right] \mathbf{u}_{n+1}^* &= \left[\mathbb{I} - \frac{1}{2}\tau \text{Adv}(\mathbf{u}_n) + \frac{1}{2}\tau\nu\Delta \right] \mathbf{u}_n \\ \Delta p_{n+1} &= \nabla \cdot \mathbf{u}_{n+1}^* \\ \mathbf{u}_{n+1} &= \mathbf{u}_{n+1}^* - \nabla p_{n+1} \end{aligned}$$

If the velocity \mathbf{u} is treated as a two-form (flux $\star \mathbf{u}^b$) and the pressure p as a three-form (integral per cell), the advection terms Adv can be discretized to produce the traditional Harlow-Welsh scheme [92] as discussed in [93]. Alas, this type of simplistic integration scheme is only relevant for low Reynolds numbers. For more involved non-linear integrator schemes, our construction of operator-adapted basis functions can still pay off significantly: one can linearize the integrator, use only a reduced basis of divergence-free velocity fields (in order to bypass the pressure projection altogether), and construct operator-adapted basis functions at the current time step in almost linear time; from these basis functions, the non-linear update rule can be solved efficiently by picking only the most relevant degrees of freedom as discussed earlier, which cuts down on the computational time of the simulation significantly (see Alg. 4). In addition, other simple heuristics can be applied; for instance, one could omit wavelets that are far enough from certain areas of interest in physical space. Finally, note that fluid simulations generally require high resolutions to accurately resolve interactions between small and large scale structures. They often benefit from subgrid models, aimed at modeling small scale effects without explicitly resolving them, thus reducing the dimensionality of the system. Our operator-adapted bases also provide a framework to directly modulate particular frequency subbands (with localization in physical space if needed) by appropriately varying the corresponding wavelet coefficients in order, for instance, to reproduce Kolmogorov's law of energy cascading. The same idea can also be applied in the more general context of fluid or smoke simulations in graphics by locally exciting or repressing certain frequencies of a simulation to produce a visually richer and/or desired look: the tools afforded by our operator-adapted basis functions and wavelets can help with a large spectrum of numerical tasks. A thorough investigation of operator-adapted wavelets for fluid dynamics is left to future work.

6. Conclusions

In this paper, we have introduced an operator-adapted multiresolution analysis for differential forms. Building over the recent notion of “gamblets”, we presented an efficient fine-to-coarse construction, in log-linear time, of a hierarchy of basis functions for differential forms and associated wavelets tailored to a given differential operator. Our numerical procedure involves only linear algebra, and can thus be easily implemented. This construction can be leveraged in finite element analysis for model reduction, numerical homogenization, or simply to provide a Galerkin discretization of an operator such that the resulting stiffness matrix becomes block-diagonal, with uniformly well-conditioned and sparse blocks. Our method also leads to what we believe is the first constructive approach to obtain vector-valued wavelets that block-diagonalize a given continuous, linear, bijective, positive-definite and self-adjoint operator, where vector fields are expressed through their covariant counterparts, i.e., one-forms. Our construction is very general: it applies to various types of computational grids, different smoothness orders of refinable basis functions, many strategies of sparsification, etc. Depending on the precise application that one targets, this flexibility is beneficial as it can accommodate a variety of specific numerical requirements.

Future work. Besides the large number of potential practical applications that this efficient construction of operator-adapted wavelets could impact, a few obvious theoretical developments are worth investigating. For instance, our construction assumes that the operator \mathcal{L} is self-adjoint and positive-definite—otherwise, we use $\mathcal{L}^T \mathcal{L}$. This fallback symmetrized operator implies a doubling of the condition number, which may not be optimal. Also, our treatment of arbitrary domain shapes within the computational grid is currently restricted to a low-order approximation of the Hodge star on the finest level. Finding higher-order spatial homogenizations of the Hodge star may be valuable to allow for broader applications. Finally, while we described our approach on “flat” domains, the notion of forms carries over naturally to non-flat domains (defined for instance through subdivision of cell complexes) as well, and a proper numerical treatment of the resulting Riemannian metric induced by the embedding space can be efficiently achieved via a push-pull quadrature of the Hodge star at the finest level [84]. This should allow an even more general treatment of operator-adapted wavelets. In particular, the construction of scale spaces (multiresolution descriptions) of geometric shapes based on the Laplace-Beltrami operator for instance could also have important applications for geometric encoding [102], without requiring a global parameterization.

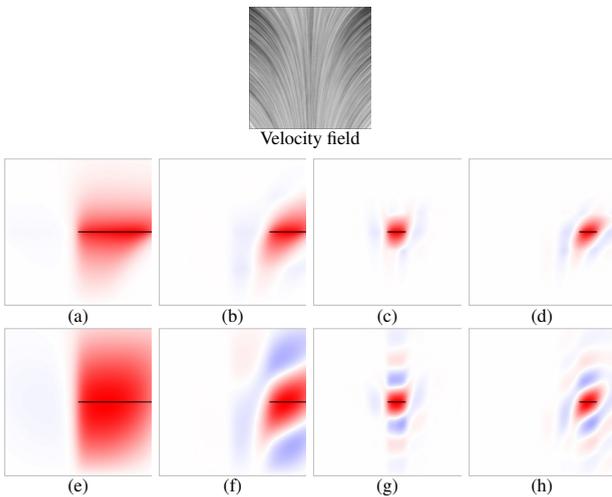


Fig. 19: **Advection-diffusion adapted basis functions.** We visualize the vertical components of 1-form basis functions φ_i^k adapted to an (upwind-evaluated) advection-diffusion operator and associated to various horizontal edges (in thick black) from 3 resolution levels using a linear red-to-blue color ramp: using Dirac-Whitney refinement at (a) level $k = 1$, (b) level $k = 2$, (c)-(d) level $k = 3$; using the original Whitney refinement at (e) level $k = 1$, (f) level $k = 2$, (g)-(h) level $k = 3$. The advecting velocity field (left) is reflected in the shapes of the adapted edge basis functions.

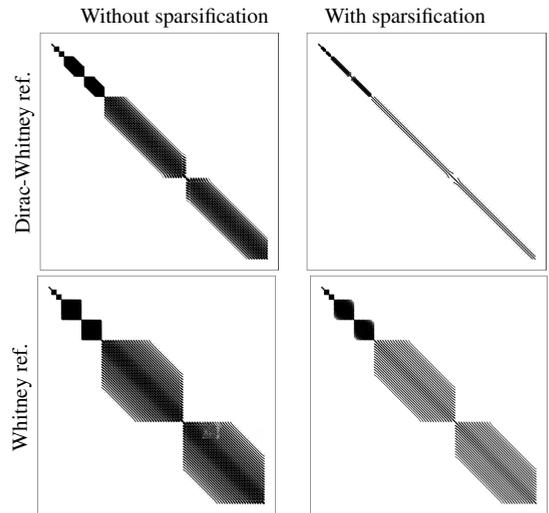


Fig. 20: **Sparsity.** We visualize the sparsity pattern of the block-diagonal stiffness matrix $\text{diag}[\mathbb{A}^1, \mathbb{B}^1, \dots, \mathbb{B}^6]$ obtained via Alg. 1 when a simple linear solver with threshold-based ($< 1e^{-15}$) pruning (left) vs. fast localized solve (right) is used, for a 1-form Laplace operator adapted decomposition using Dirac-Whitney (top) and original Whitney (bottom) refinement rules. Fast localized solves decrease sparsity (proportion of non-zeros) from 6.58% to 0.51% for Dirac-Whitney, and from 10.96% to 3.63% for original Whitney.

Appendix A. Properties of the bottom-up construction

In this first appendix, we briefly review the properties of the (scalar-valued) gamblets construction derived in [32], as they apply to our extension as well. Readers are referred to [32] for rigorous proofs: we only provide a summary of the properties, along with the key conditions required.

Bounded condition numbers. Let $\|\cdot\|_{\mathcal{L}^{-1}}$ be the norm of H^* defined as the dual to the energy norm $\|\cdot\|_{\mathcal{L}}$. The condition numbers of stiffness matrices \mathbb{B}^k and \mathbb{A}^k are uniformly bounded, provided the spans of the test functions $\{\mathcal{V}^k\}_{k=1}^q$ are regular and weakly aligned with the eigensubspaces of \mathcal{L}^{-1} in the following sense: there exists a constant c_1 , such that

- the coarsest test functions can capture the eigenspaces of \mathcal{L} , i.e.,

$$\sup_{\mathbf{x}, \mathbf{y} \in \mathbb{R}^{n_1}; |\mathbf{x}|=|\mathbf{y}|=1} \frac{\|\sum_{i=1}^{n_1} \mathbf{x}_i \varphi_i^1\|_{\mathcal{L}^{-1}}}{\|\sum_{i=1}^{n_1} \mathbf{y}_i \varphi_i^1\|_{\mathcal{L}^{-1}}} \leq c_1;$$

- the condition numbers of the wavelet refinement matrices are bounded for $k = 1, \dots, q-1$:

$$\text{cond}(\mathbf{W}^k \mathbf{W}^{k,T}) \leq c_1;$$

- the refinement matrices of test functions are bounded for $k = 1, \dots, q-1$:

$$\|\mathbf{C}^k\|_2 \leq c_1;$$

- eigensubspaces of \mathcal{L} and test functions must be weakly aligned for $k = 2, \dots, q$, i.e.,

$$\sup_{\mathbf{x} \in \text{Ker } \mathbf{C}^{k-1}, |\mathbf{x}|=1} \inf_{\mathbf{y} \in \mathbb{R}^{n_{k-1}}} \frac{\left\| \sum_{i=1}^{n_k} x_i \varphi_i^k - \sum_{j=1}^{n_{k-1}} y_j \varphi_j^{k-1} \right\|_{\mathcal{L}^{-1}}^2}{\inf_{\mathbf{z} \in \mathbb{R}^{n_k}, |\mathbf{z}|=1} \left\| \sum_{i=1}^{n_k} z_i \varphi_i^k \right\|_{\mathcal{L}^{-1}}^2} \leq c_1.$$

This inequality can also be interpreted as a bound on the relative gap between information that is lost vs. the one that is propagated during transition from fine to coarse levels. Note that slightly stronger versions of all these conditions can provide bounds for the *minimum* and *maximum* eigenvalues of matrices \mathbb{A}^k and \mathbb{B}^k , revealing that the eigenranges of \mathbb{B}^k correspond to k frequency subbands of the original stiffness matrix \mathbb{A}^q corresponding to the finest level.

Exponential decay. The fast decay of our operator-adapted basis functions on a given resolution level k follows from the localization properties of both the underlying operator and of the test functions. For each $i = 1..n^k$, define $\tau_i^k \in \Omega$ to be a small convex region including the support of the test function φ_i^k , containing a ball of radius h_k , and being contained in a ball of radius δh_k for some constants $\delta, h_k \in [0, 1]$. Let Ω_i^k be a small neighborhood of τ_i^k , such that $\Omega = \cup_{i=1}^{n^k} \Omega_i^k$, and the distance between τ_i^k and complement of Ω_i^k is between δh and h . Denoting the standard Sobolev norm (with appropriate boundary conditions) on $H(\Omega)$ by $\|\cdot\|_{H(\Omega)}$ and its dual norm on $H^*(\Omega)$ by $\|\cdot\|_{H^*(\Omega)}$, the key (sufficient) condition required for exponential decay of operator-adapted basis is the existence of two constants $0 < c_{\min} \leq c_{\max} < \infty$ that satisfy the following frame inequality for all $v \in H^*(\Omega)$:

$$c_{\min} \inf_{\varphi \in \mathcal{V}^k} \|v - \varphi\|_{H^*(\Omega)}^2 \leq \sum_{i=1}^{n_k} \inf_{\varphi \in \mathcal{V}^k} \|v - \varphi\|_{H^*(\Omega_i^k)}^2 \leq c_{\max} \inf_{\varphi \in \mathcal{V}^k} \|v - \varphi\|_{H^*(\Omega)}^2$$

The exponent of the decay rate depends on c_{\min} , c_{\max} and the locality of the operator \mathcal{L} . This condition, describing the localization of test functions in the dual norm, is in fact implied by three simple and natural inequalities (for some finite constant $c_2 > 0$):

- *Poincaré inequalities:* any φ^\perp from the L^2 -complement of test functions $\mathcal{V}^{k,\perp} = \{v \in H : \langle v, \varphi_i^k \rangle_{L^2} = 0 \forall \varphi_i^k \in \mathcal{V}^k\}$ has bounded derivatives:

$$\|D^t \varphi^\perp\|_{L^2(\Omega)} \leq c_2 h_k^{s-t} \|\varphi^\perp\|_{H_0^s(\Omega)} \quad \forall t \in \{0, \dots, s\};$$

- *Frame inequalities (boundedness of test functions):*

$$\sum_{i=1}^{n_k} \langle \varphi_i^k, f \rangle_{L^2}^2 \leq c_2 \sum_{t=0}^s h_k^{2t} \|D^t f\|_{L^2(\Omega)}^2 \quad \forall f \in H_0^s(\Omega)$$

- *Inverse Poincaré inequalities:*

$$h_k^{2s} \leq c_2 \|\varphi_i^k\|_{H^s(\tau_i)}^2 \quad \forall i \in \{1, \dots, n_k\}.$$

Note that these conditions, obtained in [32, 30], provide a generalization of [47, 28, 103].

Homogenization property. The operator-adapted basis functions on coarser levels carry over information from finer scales, nicely approximating the operator eigensubspaces of corresponding frequencies. Indeed, their resulting shape better captures the operator on a given resolution level than generic polynomial finite-element basis functions, with a bound proven analytically in [32] for several cases. More concretely, the solution u to Eq. (1) and its operator-adapted level- k FEM approximation u^k (Eq. (19)), assembled using k coarsest resolution levels, satisfy the following inequality for some constant $c > 0$:

$$\frac{\|u - u^k\|_{\mathcal{L}}}{\|g\|_{L^2}} \leq ch_k^s \quad (\text{A.1})$$

where the right-hand side is independent of \mathcal{L} , where h_k characterizes the radius of support of a test functions of level k . This last inequality implies that most of the ‘‘energy’’ concentrates on coarser resolutions and serves as one of the key reasons behind the homogenization properties of operator-adapted multiresolution analysis and its ability to use only a few discretization levels to get good approximate solutions: it guarantees that omitting (some or all) information from higher resolution levels does not degrade accuracy too much.

Appendix B. Variational definition

We briefly summarize in this appendix the original motivation of the operator-adapted decomposition proposed in [28, 29]. Since it matches our construction in the case of scalar-valued basis functions, it is interesting to understand the functional approximation roots of their approach to contrast it with our finite element perspective.

Optimal choice of basis functions. Consider the solution u to Eq. (1), and a finite set of measurements $m_i = \langle u, \varphi_i^k \rangle_{L^2}$ on a fixed level k obtained by integrating u against a set of test functions $\{\varphi_i^k\}_{i=1}^{n_k}$. In a context of functional approximation, a natural thing to ask is: what is the optimal selection of ‘‘adapted’’ basis functions φ_i^k , such that the approximation error between u and $\tilde{u}^k = \sum_{i=1}^{n_k} m_i \varphi_i^k$ is minimized, for any choice of u .

Game theoretical insight. This functional approximation problem can be formalized as a *zero-sum game* between two players as follows: player I chooses an arbitrary function $u^I \in H$ from the solution space of Eq. (1); player II constructs an approximation u^{II} of u^I from partial information, only having access to the measurements $\{m_i\}_{i=1}^{n_k}$ of player I’s function; the utility that player I aims to maximize and player II to minimize is given by:

$$V(u^I, u^{II}) = \frac{\|u^I - u^{II}\|_{\mathcal{L}}}{\|u^I\|_{\mathcal{L}}}. \quad (\text{B.1})$$

This game has a solution in mixed strategies, with mixed optimal strategy for player I, and *pure* optimal strategy for player II [28]. In particular, optimal strategy for player I consists in drawing u^I at random from a weak Gaussian distribution ζ of covariance \mathcal{L}^{-1} , while a minimax strategy for player II reduces to computing deterministic function u^{II} as the conditional expectation of ζ given the measurements $\{m_i\}_{i=1}^{n_k}$.

Variational formulation. The optimal basis functions $\{\varphi_i^k\}_{i=1}^{n_k}$ corresponding to the minimax strategy are, in fact, solutions of following variational problem:

$$\varphi_i^k = \underset{\phi \in H}{\operatorname{argmin}} \|\phi\|_{\mathcal{L}}^2 \quad \text{s.t.} \quad \langle \phi, \varphi_j^k \rangle_{L^2} = \delta_{ij} \quad \text{for } j = 1 \dots n_k \quad (\text{B.2})$$

Using this formulation, [28] shows that since the test functions were chosen to be refinable, the optimal basis functions are refinable as well:

$$\varphi_i^k = \sum_{j=1}^{n_{k+1}} \mathbb{C}_{ij}^k \varphi_j^{k+1},$$

where entries of the refinement matrix \mathbb{C}_{ij}^k correspond to conditional expectations of finer test measurements $\langle \varphi_j^{k+1}, \zeta \rangle_{L^2}$ given a single non-zero coarser level measurement $\langle \varphi_l^k, \zeta \rangle_{L^2} = \delta_{il}$ for $l = 1 \dots n_k$. To find the explicit expression of the refinement matrix, we can rewrite the problem (B.2) in matrix form, and using earlier notation $\mathbb{A}_{ij}^k := \langle \varphi_i^k, \varphi_j^k \rangle_{\mathcal{L}}$, we obtain:

$$\mathbb{C}^k = \underset{\mathbb{X} \in \mathbb{R}^{n_k \times n_{k+1}}}{\operatorname{argmin}} \operatorname{Tr} \left[\mathbb{X} \mathbb{A}^{k+1} \mathbb{X}^T \right] \quad \text{s.t.} \quad \mathbb{X} \mathbb{C}^{k,T} = \mathbb{I}_{n_k}$$

The solution to this constrained minimization is precisely Eq. (15), which demonstrates the equivalence of variational formulation of Eq. (B.1) to our axiomatic construction (Sec. 3) and implies that resulting operator-orthogonal basis functions are optimal in the game-theoretical functional-approximation sense of Eq. (B.1).

They are also optimal in the Galerkin sense on any resolution level k [28]: for any $u \in H$ and its approximation $u^k = \sum_{i=1}^{n_k} m_i \varphi_i^k$, we have

$$\|u - u^k\|_{\mathcal{L}} = \inf_{v \in \mathcal{V}^k} \|u - v\|_{\mathcal{L}}$$

In other words, u^k is the \mathcal{L} -orthogonal projection of u onto \mathcal{V}^k .

Acknowledgments

H. Owhadi gratefully acknowledges support by the Air Force Office of Scientific Research and the DARPA EQUiPS Program under award number FA9550-16-1-0054 (Computational Information Games), as well as the Air Force Office of Scientific Research under award number FA9550-18-1-0271 (Games for Computation and Learning). M. Desbrun gratefully acknowledges partial support from Pixar Animation Studios, and thanks Jiong Chen and Yiying Tong for various discussions on the topic.

References

- [1] Y. Meyer, *Wavelets and Operators*, volume 1, Cambridge university press, 1995.
- [2] G. Beylkin, R. Coifman, V. Rokhlin, Fast wavelet transforms and numerical algorithms I, *Communications on Pure and Applied Mathematics* 44 (1991) 141–183.
- [3] E. Bacry, S. Mallat, G. Papanicolaou, A wavelet space-time adaptive scheme for partial differential equations, in: *Progress in Wavelet Analysis and Applications*, Frontières, 1993, pp. 677–682.
- [4] S. Dahlke, I. Weinreich, Wavelet-Galerkin methods: an adapted biorthogonal wavelet basis, *Constr. Approx.* 9 (1993) 237–262.
- [5] S. Dahlke, I. Weinreich, Wavelet bases adapted to pseudodifferential operators, *Appl. Comput. Harmon. Anal.* 1 (1994) 267–283.
- [6] S. Bertoluzza, Y. Maday, J.-C. Ravel, A dynamically adaptive wavelet method for solving partial differential equations, *Comput. Methods Appl. Mech. Engrg.* 116 (1994) 293–299. ICOSAHOM'92 (Montpellier, 1992).
- [7] B. Engquist, S. Osher, S. Zhong, Fast wavelet based algorithms for linear evolution equations, *SIAM Journal on Scientific Computing* 15 (1994) 755–775.
- [8] O. V. Vasilyev, S. Paolucci, A dynamically adaptive multilevel wavelet collocation method for solving partial differential equations in a finite domain, *J. Comput. Phys.* 125 (1996) 498–512.
- [9] G. Chiavassa, J. Liandrat, A fully adaptive wavelet algorithm for parabolic partial differential equations, *Appl. Numer. Math.* 36 (2001) 333–358.
- [10] A. Cohen, W. Dahmen, R. DeVore, Adaptive wavelet methods for elliptic operator equations: convergence rates, *Math. Comp.* 70 (2001) 27–75.
- [11] W. Dahmen, A. Kunoth, Adaptive wavelet methods for linear-quadratic elliptic control problems: convergence rates, *SIAM J. Control Optim.* 43 (2005) 1640–1675.
- [12] R. Stevenson, Adaptive wavelet methods for solving operator equations: an overview, in: *Multiscale, Nonlinear and Adaptive Approximation*, Springer, 2009, pp. 543–597.
- [13] T. Gantumur, R. Stevenson, Computation of differential operators in wavelet coordinates, *Mathematics of Computation* 75 (2006) 697–709.
- [14] G. Beylkin, On multiresolution methods in numerical analysis, *Doc. Math., Extra 3* (1998) 481–490.
- [15] D. Gines, G. Beylkin, J. Dunn, LU factorization of non-standard forms and direct multiresolution solvers, *Applied and Computational Harmonic Analysis* 5 (1998) 156–201.
- [16] W. Dahmen, H. Harbrecht, R. Schneider, Compression techniques for boundary integral equations—asymptotically optimal complexity estimates, *SIAM Journal on Numerical Analysis* 43 (2006) 2251–2271.
- [17] J. Fröhlich, K. Schneider, An adaptive wavelet Galerkin algorithm for one- and two-dimensional flame computations, *European J. Mech. B Fluids* 13 (1994) 439–471.
- [18] B. Sendov, Adapted multiresolution analysis and wavelets, in: *Functions, series, operators*, János Bolyai Math. Soc., Budapest, 2002, pp. 23–38.
- [19] A. Cohen, I. Daubechies, J.-C. Feauveau, Biorthogonal bases of compactly supported wavelets, *Communications on Pure and Applied Mathematics* 45 (1992) 485–560.
- [20] W. Sweldens, The lifting scheme: A construction of second generation wavelets, *SIAM Journal on Mathematical Analysis* 29 (1998) 511–546.
- [21] J. M. Carnicer, W. Dahmen, J. M. Peña, Local decomposition of refinable spaces and wavelets, *Applied and Computational Harmonic Analysis* 3 (1996) 127–153.
- [22] P. S. Vassilevski, J. Wang, Stabilizing the hierarchical basis by approximate wavelets. I. Theory, *Numer. Linear Algebra Appl.* 4 (1997) 103–126.
- [23] M. Lounsbery, T. D. DeRose, J. Warren, Multiresolution analysis for surfaces of arbitrary topological type, *ACM Transactions on Graphics (TOG)* 16 (1997) 34–73.
- [24] R. Sudarshan, *Operator-adapted Finite Element Wavelets: theory and applications to a posteriori error estimation and adaptive computational modeling*, Ph.D. thesis, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, 2005.
- [25] W. Kohn, Analytic properties of Bloch waves and Wannier functions, *Physical Review* 115 (1959) 809.
- [26] G. H. Wannier, Dynamics of band electrons in electric and magnetic fields, *Reviews of Modern Physics* 34 (1962) 645.
- [27] N. Marzari, D. Vanderbilt, Maximally localized generalized Wannier functions for composite energy bands, *Physical Review B* 56 (1997) 12847.
- [28] H. Owhadi, Multigrid with rough coefficients and multiresolution operator decomposition from hierarchical information games, *SIAM Review* 59 (2017) 99–149.

- [29] H. Owjadi, L. Zhang, Gamblets for opening the complexity-bottleneck of implicit schemes for hyperbolic and parabolic ODEs/PDEs with rough coefficients, *Journal of Computational Physics* 347 (2017) 99–128.
- [30] H. Owjadi, C. Scovel, Universal scalable robust solvers from computational information games and fast eigenspace adapted multiresolution analysis, arXiv:1703.10761 (2017).
- [31] F. Schäfer, T. J. Sullivan, H. Owjadi, Compression, inversion, and approximate PCA of dense kernel matrices at near-linear computational complexity, *Computing Research Repository* abs/1706.02205 (2017).
- [32] H. Owjadi, C. Scovel, Operator adapted wavelets, fast solvers, and numerical homogenization: from a game theoretic approach to numerical approximation and algorithm design, *Cambridge Monographs on Appl. Comput. Math.*, Cambridge University Press, 2019.
- [33] É. Cartan, *Les systèmes différentiels extérieurs et leurs applications géométriques* (1945).
- [34] D. N. Arnold, R. S. Falk, R. Winther, Finite element exterior calculus: from Hodge theory to numerical stability, *Bull. Amer. Math. Soc. (N.S.)* 47 (2010) 281–354.
- [35] K. Amaratunga, J. R. Williams, S. Qian, J. Weiss, Wavelet-Galerkin solutions for one-dimensional partial differential equations, *International Journal for Numerical Methods in Engineering* 37 (1994) 2703–2716.
- [36] M. Brewster, G. Beylkin, A multiresolution strategy for numerical homogenization, *Applied and Computational Harmonic Analysis* 2 (1995) 327–349.
- [37] M. Dorobantu, B. Engquist, Wavelet-based numerical homogenization, *SIAM Journal on Numerical Analysis* 35 (1998) 540–559.
- [38] G. Beylkin, N. Coult, A multiresolution strategy for reduction of elliptic PDEs and eigenvalue problems, *Applied and Computational Harmonic Analysis* 5 (1998) 129–155.
- [39] G. R. Yoo, H. Owjadi, De-noising by thresholding operator adapted wavelets, arXiv preprint arXiv:1805.10736 (2018).
- [40] H. Xie, L. Zhang, H. Owjadi, Fast eigenpairs computation with operator adapted wavelets and hierarchical subspace correction, arXiv preprint arXiv:1806.00565 (2018).
- [41] A. Bossavit, Whitney forms: a class of finite elements for three-dimensional computations in electromagnetism, *IEEE Proceedings A* 135 (1988) 493–500.
- [42] D. N. Arnold, Spaces of finite element differential forms, in: U. Gianazza, F. Brezzi, P. Colli Franzone, G. Gilardi (Eds.), *Analysis and Numerics of Partial Differential Equations*, Springer, 2013, pp. 117–140.
- [43] M. Desbrun, E. Kanso, Y. Tong, Discrete differential forms for computational modeling, in: A. Bobenko, et al. (Eds.), *Discrete Differential Geometry*, Birkhäuser Basel, 2008, pp. 287–324.
- [44] C. A. Micchelli, T. J. Rivlin, A survey of optimal recovery, in: *Optimal Estimation in Approximation Theory*, Springer, 1977, pp. 1–54.
- [45] P. S. Vassilevski, General constrained energy minimization interpolation mappings for AMG, *SIAM J. Sci. Comput.* 32 (2010) 1–13.
- [46] T. J. R. Hughes, G. R. Feijóo, L. Mazzei, J.-B. Quinicy, The variational multiscale method—a paradigm for computational mechanics, *Computer Methods in Applied Mechanics and Engineering* 166 (1998) 3–24.
- [47] A. Målqvist, D. Peterseim, Localization of elliptic multiscale problems, *Mathematics of Computation* 83 (2014) 2583–2603.
- [48] R. L. Harder, R. N. Desmarais, Interpolation using surface splines, *J. Aircraft* 9 (1972) 189–191.
- [49] J. Duchon, Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces, *Rev. Française Automat. Informat. Recherche Operationnelle Ser. RAIRO Analyse Numerique* 10 (1976) 5–12.
- [50] H. Owjadi, L. Zhang, L. Berlyand, Polyharmonic homogenization, rough polyharmonic splines and sparse super-localization, *ESAIM: Mathematical Modelling and Numerical Analysis* 48 (2014) 517–552.
- [51] J. Chen, H. Bao, T. Wang, M. Desbrun, J. Huang, Numerical coarsening using discontinuous shape functions, *ACM Trans. Graph.* 37 (2018) Art. 117.
- [52] P.-A. Raviart, J. M. Thomas, A mixed finite element method for 2nd order elliptic problems 606 (1977) 292–315.
- [53] J.-C. Nédélec, Mixed finite elements in \mathbf{R}^3 , *Numer. Math.* 35 (1980) 315–341.
- [54] A. Bossavit, Mixed finite elements and the complex of Whitney forms, *The mathematics of finite elements and applications VI* (1988) 137–144.
- [55] H. Whitney, *Geometric integration theory*, Princeton University Press, 1957.
- [56] J. R. Munkres, *Elements of Algebraic Topology*, Addison-Wesley, 1984.
- [57] D. N. Arnold, R. S. Falk, R. Winther, Finite element exterior calculus, homological techniques, and applications, *Acta Numerica* 15 (2006) 1–155.
- [58] J.-C. Nédélec, Mixed finite elements in \mathbb{R}^3 , *Numer. Math.* 35 (1980) 315–341.
- [59] A. Stern, Y. Tong, M. Desbrun, J. E. Marsden, Variational integrators for Maxwell’s equations with sources, *PIERS Online* 4 (2008) 711–715.
- [60] E. Gawlik, P. Mullen, D. Pavlov, J. Marsden, M. Desbrun, Geometric, variational discretization of continuum theories, *Physica D: Nonlinear Phenomena* 240 (2011) 1724–1760.
- [61] B. Liu, G. Mason, J. Hodgson, Y. Tong, M. Desbrun, Model-reduced variational fluid simulation, *ACM Trans. Graph.* 34 (2015) Art. 244.
- [62] A. N. Hirani, K. B. Nakshatrala, J. H. Chaudhry, Numerical method for Darcy flow derived using Discrete Exterior Calculus, *International Journal for Computational Methods in Engineering Science & Mechanics* 16 (2015) 151–169.
- [63] M. S. Mohamed, A. N. Hirani, R. Samtaney, Discrete exterior calculus discretization of incompressible Navier–Stokes equations over surface simplicial meshes, *Journal of Computational Physics* 312 (2016) 175–191.
- [64] Y. Tong, P. Alliez, D. Cohen-Steiner, M. Desbrun, Designing quadrangulations with discrete harmonic forms, in: *Symposium on Geometry Processing*, 2006, pp. 201–210.
- [65] L. J. Grady, J. R. Polimeni, *Discrete Calculus: Applied Analysis on Graphs for Computational Science*, Springer-Verlag, 2010.
- [66] A. Weil, Sur les théorèmes de de Rham, *Comment. Math. Helv.* 26 (1952) 119–145.
- [67] H. Whitney, *Geometric integration theory*, Princeton University Press, 1957.
- [68] P. M. Anselone, P.-J. Laurent, A general method for the construction of interpolating or smoothing spline-functions, *Numerische Mathematik* 12 (1968) 66–82.
- [69] I. J. Schoenberg, Splines and histograms, in: *Spline Functions and Approximation Theory*, Springer, 1973, pp. 277–327.
- [70] N. Robidoux, Polynomial histopolation, superconvergent degrees of freedom, and pseudospectral discrete Hodge operators, 2008. Unpublished, found at: people.math.sfu.ca/~nrobidou/public.html/prints/histogram/histogram.pdf.

- [71] R. Abraham, J. E. Marsden, R. Ratiu, *Manifolds, Tensor Analysis, and Applications: 2nd Edition*, Springer-Verlag, 1988.
- [72] R. Hiptmair, Higher order Whitney forms, *Progress In Electromagnetics Research* 32 (2001) 271–299.
- [73] R. Hiptmair, Finite elements in computational electromagnetism, *Acta Numerica* 11 (2002) 237–339.
- [74] F. Rapetti, A. Bossavit, Whitney forms of higher degree, *SIAM J. Numer. Anal.* 47 (2009) 2369–2386.
- [75] K. Wang, Weiwei, Y. Tong, M. Desbrun, P. Schröder, Edge subdivision schemes and the construction of smooth vector fields, *ACM Trans. Graph.* 25 (2006) 1041–1048.
- [76] S. H. Christiansen, F. Rapetti, On high order finite element spaces of differential forms, *Math. Comp.* 85 (2016) 517–548.
- [77] A. Gillette, A. Rand, C. Bajaj, Construction of scalar and vector finite element families on polygonal and polyhedral meshes, *Computational Methods in Applied Mathematics* 16 (2016) 667–683.
- [78] A. Buffa, J. Rivas, G. Sangalli, R. Vázquez, Isogeometric discrete differential forms in three dimensions, *SIAM Journal on Numerical Analysis* 49 (2011) 818–844.
- [79] W. Chen, Y. Wang, Minimal degree $h(\text{curl})$ and $h(\text{div})$ conforming finite elements on polytopal meshes, *Math. Comp.* 86 (2017) 2053–2087.
- [80] A. Gillette, T. Kloefkorn, V. Sanders, Computational Serendipity and Tensor Product Finite Element Differential Forms, *ArXiv e-prints math.NA 1806.00031* (2018).
- [81] D. Rufat, G. Mason, P. Mullen, M. Desbrun, The chain collocation method: a spectrally accurate calculus of forms, *Journal of Computational Physics* 257 (2014) 1352–1372.
- [82] M. Gerritsma, Edge functions for spectral element methods, in: *Spectral and High Order Methods for Partial Differential Equations*, Springer, 2011, pp. 199–207.
- [83] E. Grinspun, The basis refinement method, Ph.D. thesis, California Institute of Technology, Engineering and Applied Science, 2003.
- [84] F. de Goes, M. Desbrun, M. Meyer, T. DeRose, Subdivision exterior calculus for geometry processing, *ACM Trans. Graph.* 35 (2016) Art. 133.
- [85] D. Zorin, P. Schröder, Subdivision for modeling and animation, 2000. Course notes at ACM SIGGRAPH Conference.
- [86] A. Hirani, Discrete Exterior Calculus, Ph.D. thesis, Computer Science, California Institute of Technology, 2003.
- [87] L. Feng, P. Alliez, L. Busé, H. Delingette, M. Desbrun, Curved optimal Delaunay triangulation, *ACM Trans. Graph.* 37 (2018) Art. 61.
- [88] K. Wang, A subdivision approach to the construction of smooth differential forms, Ph.D. thesis, Applied Mathematics, California Institute of Technology, 2008.
- [89] J. Stam, Exact evaluation of Catmull-Clark subdivision surfaces at arbitrary parameter values, in: *Proceedings of SIGGRAPH Conference*, 1998, pp. 395–404.
- [90] P. B. Bochev, J. M. Hyman, Principles of mimetic discretizations of differential operators, in: D. N. Arnold, et al. (Eds.), *Compatible Spatial Discretizations*, Springer New York, 2006, pp. 89–119.
- [91] P. Ciarlet, *The Finite Element Method for Elliptic Problems*, SIAM, 2002.
- [92] F. H. Harlow, J. E. Welch, Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface, *Physics of Fluids* 8 (1965) 2182–2189.
- [93] D. Pavlov, P. Mullen, Y. Tong, E. Kanso, J. E. Marsden, M. Desbrun, Structure-preserving discretization of incompressible fluids, *Physica D: Nonlinear Phenomena* 240 (2011) 443–458.
- [94] Y. T. Ng, C. Min, F. Gibou, An efficient fluid-solid coupling algorithm for single-phase flows, *J. Comput. Phys.* 228 (2009) 8807–8829.
- [95] B. Cabral, L. C. Leedom, Imaging vector fields using line integral convolution, in: *Proceedings of SIGGRAPH Conference*, 1993, pp. 263–270.
- [96] D. Rufat, Licpy: Line integral convolution with python, 2018. <https://github.com/drufat/licpy>.
- [97] P. Mullen, A. McKenzie, D. Pavlov, L. Durant, Y. Tong, E. Kanso, J. E. Marsden, M. Desbrun, Discrete Lie advection of differential forms, *Foundations of Computational Mathematics* 11 (2011) 131–149.
- [98] K. Urban, On divergence-free wavelets, *Advances in Computational Mathematics* 4 (1995) 51–81.
- [99] E. Deriaz, V. Perrier, Divergence-free and curl-free wavelets in two dimensions and three dimensions: application to turbulent flows, *Journal of Turbulence* 7 (2006) N3.
- [100] I. Babuška, J. E. Osborn, Can a finite element method perform arbitrarily badly?, *Mathematics of Computation* 69 (2000) 443–462.
- [101] T. Bui-Thanh, K. Willcox, O. Ghattas, B. van Bloemen Waanders, Goal-oriented, model-constrained optimization for reduction of large-scale systems, *Journal of Computational Physics* 224 (2007) 880–896.
- [102] A. Khodakovskiy, P. Schröder, W. Sweldens, Progressive geometry compression, in: *Proceedings of SIGGRAPH Conference*, 2000, pp. 271–278.
- [103] R. Kornhuber, H. Yserentant, Numerical homogenization of elliptic multiscale problems by subspace decomposition, *Multiscale Modeling & Simulation* 14 (2016) 1017–1036.