

Material-adapted Refinable Basis Functions for Elasticity Simulation

JIONG CHEN, State Key Lab of CAD&CG, Zhejiang University
MAX BUDNINSKIY and HOUMAN OWHADI, Caltech
HUJUN BAO and JIN HUANG*, State Key Lab of CAD&CG, Zhejiang University
MATHIEU DESBRUN, Caltech

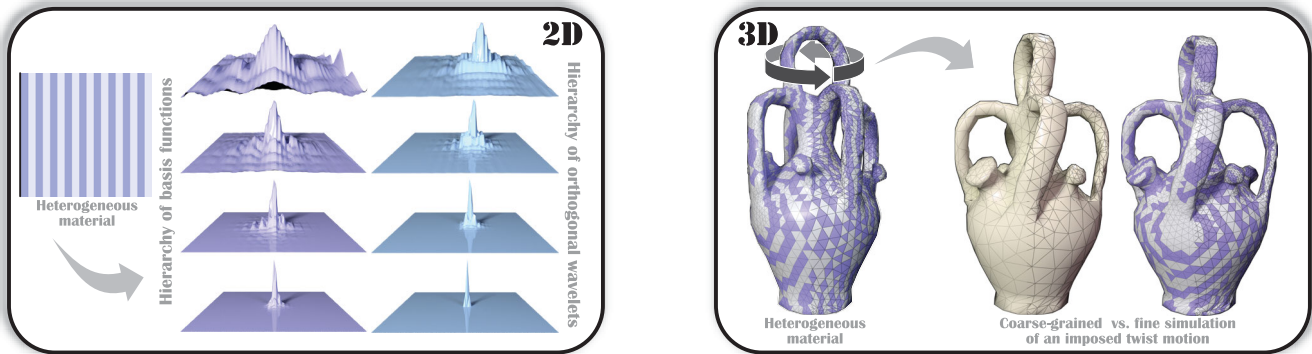


Fig. 1. **Wavelet-based coarse graining of inhomogeneous linear materials.** We introduce a new approach for numerically capturing the behavior of highly inhomogeneous linear elastic objects on coarse grids. Through a hierarchical construction of material-adapted sparse basis functions and associated wavelets (left, 2D example), we provide an efficient model reduction approach which reproduces the physical behavior of any complex material composition even at coarse scales (right). The resulting simulation involves only well conditioned sparse matrices, leading to several orders of magnitude acceleration compared to a solve based on the fine ill-conditioned linear system derived from typical piecewise polynomial basis functions.

In this paper, we introduce a hierarchical construction of material-adapted refinable basis functions and associated wavelets to offer efficient coarse-graining of linear elastic objects. While spectral methods rely on global basis functions to restrict the number of degrees of freedom, our basis functions are locally supported; yet, unlike typical polynomial basis functions, they are adapted to the material inhomogeneity of the elastic object to better capture its physical properties and behavior. In particular, they share spectral approximation properties with eigenfunctions, offering a good compromise between computational complexity and accuracy. Their construction involves only linear algebra and follows a fine-to-coarse approach, leading to a block-diagonalization of the stiffness matrix where each block corresponds to an intermediate scale space of the elastic object. Once this hierarchy has been precomputed, we can simulate an object at runtime on very coarse resolution grids and still capture the correct physical behavior, with orders of magnitude speedup compared to a fine simulation. We show on a variety of heterogeneous materials that our approach outperforms all previous coarse-graining methods for elasticity.

*Corresponding author: hj@cad.zju.edu.cn.

Authors' addresses: J. Chen, H. Bao, J. Huang, State Key Lab of CAD&CG, Zhejiang University, Hangzhou, China; M. Budninskiy, H. Owhadi, M. Desbrun, Computing + Mathematical Sciences, California Institute of Technology, Pasadena, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.

0730-0301/2019/11-ART161 \$15.00

<https://doi.org/10.1145/3355089.3356567>

CCS Concepts: • **Computing methodologies** → **Physical simulation**.

Additional Key Words and Phrases: Numerical coarsening, deformable body simulation, material-adapted basis functions, operator-adapted wavelets.

ACM Reference Format:

Jiong Chen, Max Budninskiy, Houman Owhadi, Hujun Bao, Jin Huang, and Mathieu Desbrun. 2019. Material-adapted Refinable Basis Functions for Elasticity Simulation. *ACM Trans. Graph.* 38, 6, Article 161 (November 2019), 15 pages. <https://doi.org/10.1145/3355089.3356567>

1 INTRODUCTION

Inhomogeneous elastic bodies exhibit a wide range of space and time scales, with rapid vibrations in stiff parts and slower, larger deformation in softer regions. Efficiently capturing this wide range of local behaviors numerically is challenging: using an overly fine spatial discretization to represent the inhomogeneous structure of the medium faithfully is prohibitively costly, while simply using a coarse mesh with low-order polynomial finite elements to simulate such complex materials can perform arbitrarily badly [Babuška and Osborn 2000]. As the demand for ever more complex simulation increases (with applications ranging from realtime surgical training to rapid prototyping of metamaterials), the need for scalable coarse-graining methods whose computational costs grow reasonably slowly with the structural complexity of the simulated object has become paramount.

Coarsening (or homogenization), i.e., finding an efficient method that best approximates on a coarse computational grid the behavior of the original higher complexity material, is in fact a key issue in




Fig. 2. **Homogenization effects per level.** For a fine 16×16 bi-material elastic composite (the light material being softer) hung from its fixed top, an extremely coarse 2×2 simulation using a Dirac refinement captures the proper behavior, but part of the deformation is only captured by the wavelets at the same level—the other levels carrying no discernible additional frequencies. The bilinear refinement provides better coarse-graining, since the coarsest 2×2 level already captures the complete deformation.

computational science for increased efficiency and scalability. Indeed, dynamical models with highly heterogeneous materials typically lead to ill-conditioned systems of equations, and even multigrid methods (which have been used successfully for shells [Jeon et al. 2013; Tamstorf et al. 2015] and volumes [Zhu et al. 2010; McAdams et al. 2011; Dick et al. 2011]) can only handle mostly homogeneous problems: their convergence rate can be severely affected by the lack of regularity in the material coefficients [Alcouffe et al. 1981; Yavneh 2006; Brandt et al. 2011]. This paper discusses a new homogenization method for vector-valued partial differential equations (PDEs) applicable to highly heterogeneous elastic bodies, based on the design of hierarchical material-adapted basis functions (and their associated wavelets) to allow for fast, yet accurate simulation.

1.1 Previous work

A large variety of approaches have been proposed to address efficiency and scalability in the simulation of deformable bodies. Adaptive simulation through spatial adaptation [Narain et al. 2012], wavelets [Grinspun et al. 2002] or even space and time and adaptive sampling [Debunne et al. 2001] can counteract the increased visual rigidity and artifacts that coarse elements with low-order basis functions produce, using refinements (in the number of elements and/or their polynomial orders) based on the local amount of deformation to distribute degrees of freedom where they are most needed. However, inhomogeneous materials often require a high resolution over most of the domain, rendering adaptivity unable to improve efficiency. Another strategy is to intentionally limit the space of possible deformations, through some form of modal analysis [Pentland and Williams 1989; Krysl et al. 2001; Hauser et al. 2003; Barbič and James 2005; Treuille et al. 2006; Liu et al. 2015] or even keyframing and interpolation [Martin et al. 2011]. However, the resulting algorithms do not scale well as they are typically high in runtime complexity, preventing the capture of fine details at low computational cost. Moreover, their degrees of freedom are associated with global deformations, not with spatial locations, fundamentally complicating the treatment of contacts and collisions— but this idea

has been particularly fruitful in realtime motion editing [Barbič and Popović 2008; Huang et al. 2011; Hildebrandt et al. 2012; Li et al. 2014]. Another family of approaches, called *numerical coarsening* (or *homogenization*), propose to find constitutive parameters for each coarse element that best approximate the behavior of the original heterogeneous medium [Kharevych et al. 2009; Panetta et al. 2015; Chen et al. 2015, 2017], or to construct local basis functions that are better able to capture the physical behavior of the model [Nesme et al. 2009; Torres et al. 2014; Chen et al. 2018]. These methods are particularly convenient in practice as they can directly be integrated within a conventional Galerkin simulation code.

Homogenization has, thus far, demonstrated the best compromise between efficiency and accuracy in the treatment of inhomogeneous materials; yet, current methods have a number of limitations. Element-wise coarsening of basis functions [Nesme et al. 2009; Torres et al. 2014] is often limited in its ability to approximate the proper dynamics for highly inhomogeneous materials as the junctions between elements are considered constrained, adding undue rigidity. Chen et al. [2018] thus proposed to use matrix-based and non-conforming basis functions that capture some key global deformations exactly. However, this most recent approach is not without significant shortcomings as well. First, boundary conditions are entirely ignored in the coarsening process and only enforced numerically at runtime, potentially creating significant visual artifacts (see Fig. 11); the amount of coarsening is also effectively limited as the construction of their bases require solving a large amount of local optimizations, which can only be amortized through parallelization; moreover, the actual dynamics (potentially involving damping) is not accounted for in the coarsening; finally, the authors only considered a two-scale homogenization, preventing the computational flexibility afforded by adaptive methods.

1.2 Outline

In this paper, we introduce a novel approach to the design of nested functional spaces that are adapted to the material properties of a given inhomogeneous elastic object. We will argue that these basis functions should be localized in both space *and* eigenspace, and should also respect fundamental invariances (translation and rotation) to be most relevant for simulation, and will introduce a hierarchical construction of such material-adapted basis functions, along with their associated scale-orthogonal wavelets. With the resulting elasticity-induced multiresolution decomposition of a fine functional space, we will show that single-scale or adaptive coarse-graining, augmented with a rotation correction [Huang et al. 2011] to account for geometric non-linearity, can be achieved more efficiently and with better coarse-grained results than previous methods.

2 MOTIVATIONS AND FOUNDATIONS

We now present the core concepts behind our contribution. For the sake of generality, we assume in this section that one seeks to compute a Galerkin discretization of a given linear, symmetric and positive definite operator \mathcal{L} , either in the context of solving a static problem (finding a finite-dimensional scalar field u such that $\mathcal{L}u = g$ for given forcing terms g), or for dynamic simulation (e.g., performing an integration time step of the form $u_{t+dt} = \mathcal{L}u_t$). The operator is further assumed to derive from a highly heterogeneous

and strongly anisotropic medium, rendering its solution space *multiscale*—i.e., involving a wide range of spatial scales and time scales. In this case, coarse piecewise-polynomial basis functions lead to severe inaccuracies [Babuška and Osborn 2000]. We discuss various possibilities in the choice of basis functions to allow for a reliable numerical method, even on coarse grids.

2.1 Local vs. global basis functions

Pros and cons of global support. The visual importance of the fundamental frequencies of an elastic object was noted early on in computer animation [Pentland and Williams 1989]. In essence, the eigenfunctions of the linear operator one wishes to simulate are indeed seemingly the best basis functions to rely on if one wishes to rewrite the dynamics as a function of fewer variables: keeping only a low number of degrees of freedom will still capture the most important frequencies. Since these basis functions diagonalize the operator, they can be handled individually as they are truly representing different scales with respect to this operator. In practice, however, the global nature of these eigenbases is also their main shortcoming. Having no associated spatial degree of freedom to manipulate spells trouble: spatial collision handling cannot be resolved directly in eigenspace, and any change of boundary conditions requires a new eigendecomposition. These issues are also present in methods using other non-local basis functions [Treuille et al. 2006; Martin et al. 2011; De Witt et al. 2012; Liu et al. 2015]. Moreover, computing a sufficient number of eigenmodes for very inhomogeneous materials is numerically difficult as eigenvalues often become closely clustered, bringing ill-conditioning.

Pros and cons of local support. For the simulation of homogeneous deformable bodies, the use of low-order polynomial basis functions defined over simple spatial grids is ubiquitous in animation. Alas, keeping the same local basis functions for coarse simulations of *inhomogeneous* materials fails spectacularly. Recent works have focused on tailoring local coarse basis functions such that their use in a typical Galerkin method directly on the coarse computational grid exhibits the dynamical characteristics of the original complex material. A first approach to designing material-adapted basis functions was offered in [Nesme et al. 2009] through per-element constrained static equilibrium solves. The locality of this homogenization procedure and lack of enforcement of boundary conditions was shown to be limiting in its ability to handle complex material composition [Torres et al. 2014]. Local, non-conforming and matrix-valued basis functions that are able to exactly capture key global modes of the deformable objects [Kharevych et al. 2009] were introduced recently [Chen et al. 2018], leading to much improved results on large deformation of inhomogeneous materials. However, it remains difficult to quantify how well a local basis is adapted to a given material composition: besides their important geometric properties (such as partition of unity), the bases of [Chen et al. 2018] may not, for instance, be able to enforce boundary conditions well, and they are only guaranteed to exactly capture a small number of global deformations around equilibrium. Thus, while local methods are easy to incorporate into a typical finite element simulation pipeline, they do not necessarily capture even the most fundamental frequencies of the material [Chen et al. 2019].

2.2 Multiresolution analysis

Multiscale basis functions. As spectral and local basis functions present two very different options with their own set of benefits and shortcomings, MultiResolution Analysis (MRA [Mallat 2008]) bridges these two extremes by defining a hierarchy of *local and refinable basis functions* spanning functional spaces \mathcal{V}^k (from $k=1$ for the coarsest space to $k=q$ for the finest space) satisfying $\mathcal{V}^k \subset \mathcal{V}^{k+1}$: for a given hierarchy of grids \mathcal{M}^k , nodal basis functions $\{\varphi_i^k\}_i$ are defined as local, linear combination of finer basis functions $\{\varphi_i^{k+1}\}_i$. Consequently, the spatial support size of basis functions at level k decreases with k . Each grid level k is now able to capture a given frequency band of an arbitrary function through a series of spatial degrees of freedom. In that sense, MRA offers space- and frequency-localized basis functions. Moreover, one can additionally define *wavelets* $\{\psi_i^k\}_i$ corresponding to a basis of the complement functional space between two consecutive scales; i.e., wavelets capture the details that are *missed* when a function in \mathcal{V}^{k+1} is approximated by a function in \mathcal{V}^k . Such hierarchical basis functions and their wavelets have found countless applications as they can be used in a Galerkin framework for adaptive simulation [Stollnitz et al. 1996; Grinspun et al. 2002], offering a flexible alternative to solely-global or only-local bases discussed above.

Computational limitations. Despite their early use in numerical coarsening [Brewster and Beylkin 1995; Dorobantu and Engquist 1998] due to their attractive seamless adaptivity, wavelet-based Galerkin approaches to simulation cannot handle highly inhomogeneous materials well *either*: a wavelet-Galerkin discretization can in fact worsen the condition number of the resulting system of equations. Indeed, the stiffness matrix expressed in this MRA basis becomes dense as the wavelets are not *adapted* to the operator at play. In other words, while MRA forms a decomposition where basis functions and wavelets are L_2 -orthogonal, i.e.,

$$\forall i, j, k, \int_{\Omega} \psi_i^k \varphi_j^k dx = 0;$$

their use in the context of an operator \mathcal{L} would require, instead:

$$\forall i, j, k, \int_{\Omega} \psi_i^k \mathcal{L} \varphi_j^k dx = 0. \quad (1)$$

Only then the decomposition can become optimal for computations involving \mathcal{L} : now the wavelets are truly orthogonal *with respect to the operator* to the basis functions of their own resolution and to all other wavelets, thus offering a functional representation which *block-diagonalizes* the operator \mathcal{L} . Failure to enforce this \mathcal{L} -orthogonality has obvious consequences: since the basis functions and wavelets are not aware of the fundamental frequencies of \mathcal{L} , high frequencies of the system can be present at coarse scales, and vice-versa. This notion of a \mathcal{L} -adapted MRA, different from the traditional “signal processing” standpoint and its associated use of L_2 , was first spelled out in [Sudarshan 2005], but no practical construction of \mathcal{L} -adapted basis functions and wavelets was offered.

2.3 Space- and eigenspace-localized basis functions

Variational adaptation of bases. In a recent work, Owhadi [Owhadi 2017] introduced a variational definition of optimal operator-adapted basis functions inspired by a long history of optimal functional approximation, starting with the concept of Wannier functions [Wannier 1937], as well as the more recent notions of optimal recovery

splines [Micchelli and Rivlin 1977], polyharmonic splines [Harder and Desmarais 1972] and energy-minimizing splines [Vassilevski 2010]. Given an existing set of conventional locally-supported basis functions $\{\varphi_i\}_i$ on a grid, a set of operator-adapted basis functions $\{\varphi_i\}_i$ is defined such that each one of them satisfies:

$$\varphi_i = \arg \min_{\phi} \int_{\Omega} \phi \mathcal{L} \phi \quad \text{s.t.} \quad \int_{\Omega} \phi \varphi_j = \delta_{ij} \quad \forall j. \quad (2)$$

In other words, the adapted basis functions are as close as possible to the fundamental frequency (as they minimize the quadratic form associated with \mathcal{L}) while being forced to weakly reproduce the original basis functions. Owhadi showed that the adapted basis functions thus defined turn out to be exponentially-decaying in space, which means that a simple thresholding make them tightly localized in space around the original basis location. Yet, they are also localized in eigenspace, i.e., they stand in the span of only a few eigenvectors associated to the lowest eigenvalues. This dual localization thus strikes a balance between *spectral* and *local* bases. Note also that this constrained quadratic minimization is significantly simpler than an ℓ_1 -based minimization to induce localization [Brandt and Hildebrandt 2017], and that the local support of the resulting functions is induced by the operator, not forced to be precisely the size of a coarse element like previous approaches [Nesme et al. 2009].

Operator-adapted wavelets. Furthermore, Owhadi showed that one can construct a *nested* set of operator-adapted basis functions and their associated wavelets (dubbed “gamblets” due to their game-theoretic interpretation), providing a full \mathcal{L} -adapted MRA. Based on an existing L_2 -based hierarchy of refinable basis functions and associated wavelets defined over a hierarchy of meshes $\{\mathcal{M}_k\}_k$, he described in [Owhadi 2017] how to start at the finest level with the original finest basis functions spanning a functional space \mathcal{V}^q and, in a bottom-up fashion, decompose this finest space into a hierarchy of \mathcal{L} -adapted basis functions and wavelets with simple linear algebra and *without* having to repeatedly solve the optimization problem in Eq. (2) for each of them. The result is a new set of material-adapted basis functions (that capture the most relevant \mathcal{L} -eigenspaces) and their associated wavelets, with the following properties:

- for each level k , the material-adapted basis functions φ_i^k on the nodes of mesh \mathcal{M}^k are local, and span a functional space \mathcal{V}^k ;
- For all levels but the finest, each function φ_i^k is refinable, i.e., it is a linear combination of finer adapted basis functions φ_i^{k+1} ;
- the wavelets of level k (i.e., the details in \mathcal{V}^{k+1} not in \mathcal{V}^k) span the \mathcal{L} -orthogonal complement \mathcal{W}^k of \mathcal{V}^k in \mathcal{V}^{k+1} :

$$\mathcal{V}^{k+1} = \mathcal{V}^k \oplus_{\mathcal{L}} \mathcal{W}^k, \quad (3)$$

hence Eq. (1) is satisfied at each adapted level $\neq q$;

- the constructed hierarchy of basis functions and wavelets spans the original finest functional space:

$$\mathcal{V}^q = \mathcal{V}^1 \oplus_{\mathcal{L}} \mathcal{W}^1 \oplus_{\mathcal{L}} \dots \oplus_{\mathcal{L}} \mathcal{W}^{q-2} \oplus_{\mathcal{L}} \mathcal{W}^{q-1}, \quad (4)$$

- the condition numbers of the Galerkin stiffness matrices for \mathcal{L} using the adapted basis functions in any level k are *uniformly bounded*; similarly for the wavelet-based stiffness matrices.

Since the resulting hierarchy of adapted basis functions and wavelets block-diagonalizes the operator \mathcal{L} per level, it offers a very convenient adaptive framework for efficient model reduction of the

operator [Owhadi and Zhang 2017]. Alas, this construction is fundamentally designed for *scalar-valued* operators, and thus inappropriate for elasticity: a per-coordinate treatment of the elastic displacements would entirely fail to capture the inhomogenous and anisotropic structure of the material, an issue already pointed out in [Chen et al. 2018]. Even the recent work of [Budninskiy et al. 2019], extending gamblets to finite-element differential forms, does not allow for homogenization of elastic models.

2.4 Contributions

While the approach in [Owhadi 2017] only applies to scalar-valued PDEs, we propose in this paper to extend it to vector-valued equations and apply it to the fast simulation of deformable bodies with highly heterogeneous materials. In the process, we introduce the following contributions to the field of elasticity coarsening:

- we introduce the use of basis functions that are localized in space and eigenspace, yet *respect* infinitesimal translation- and rotation-invariance for the efficient simulation of complex materials, offering far improved coarsening compared to per-coarse-element analysis of the material used in previous methods;
- following the recent construction of [Chen et al. 2018], our basis functions are matrix-valued to make sure the 3D coordinates are spatially correlated based on the local material composition, and are stored using the finest level of resolution;
- a whole hierarchy of such basis functions is constructed to offer a multiresolution numerical treatment of deformation, naturally accounting for Dirichlet conditions;
- the use of Rotation-Strain coordinates, previously used in the context of modal analysis, is leveraged in our context to robustly address geometric non-linearity (i.e., to remove the visual artifacts brought on by the linearization of the Green strain tensor that linear elasticity is based on);
- finally, our construction can handle coarsening of both elastostatics and dynamics, i.e., the homogenization of just the elasticity operator or of a time integration step for elasticity.

Next we explain our coarsening process to construct the hierarchy of adapted basis functions for a given inhomogeneous elastic material.

3 CONSTRUCTING MATERIAL-ADAPTED HIERARCHY

We now go over our approach to create a material-adapted hierarchy of refinable basis functions and their wavelets. This construction, which turns a conventional multiresolution sequence of functional spaces into a set of basis functions adapted to a given deformable body, is achieved in a fine-to-coarse manner through linear algebra. For clarity, all material-adapted entities will be denoted with the **outline** font, so that related canonical and material-adapted quantities can be denoted similarly, but distinctly; for instance, while the basis function for node i at level k is denoted φ_i^k , the associated material-adapted basis function will be, instead, denoted ϕ_i^k .

3.1 Setup

Let \mathcal{B} be the deformable body we seek to simulate.

Mesh hierarchy. While our construction applies to *arbitrary* mesh hierarchies (as we will discuss and demonstrate in Sec. 5.1), we




Fig. 3. **Mesh hierarchy.** Our method accommodates various mesh hierarchies: homogenization can be achieved over nested or non-nested refinement levels of simplicial (top) or polyhedral (bottom) mesh elements.

assume for simplicity of exposition that we are given a hierarchy of nested regular grids $\{\mathcal{M}^k\}_{k=1..q}$ in 3D (with level $k+1$ being twice as fine as level k , see Fig.3 and 2D inset) such that the finest grid \mathcal{M}^q finely discretizes the volume of \mathcal{B} in its rest shape. We denote by \bar{x}_i^k the position of node i of grid \mathcal{M}^k . Each element of this fine discretization is given its own elastic parameters (Young's modulus, Poisson's ratio, and density) corresponding to the heterogeneous medium in \mathcal{B} . A subset of boundary faces of this fine grid are also tagged, representing the parts of the boundary $\partial\mathcal{B}$ with (fixed) Dirichlet boundary conditions. Finally, we denote by n_k the number of nodes on grid \mathcal{M}^k , and we define $\eta_k \equiv n_{k+1} - n_k$ to be the loss of degrees of freedom between level k and level $k+1$.

Initial fine discretization of \mathcal{L} . On the finest grid \mathcal{M}^q , we denote by $\{\bar{\varphi}_i\}_{i=1..n_q}$ the standard nodal-based *trilinear* basis functions. From these basis functions, the stiffness matrix \mathbf{A}^q for linear elasticity given the fine material properties is assembled as in typical finite-element Galerkin simulations—see, e.g., [Irving et al. 2006].

Solution sought. Solving elastostatics or dynamics for the body \mathcal{B} consists in computing the displacement $u_i^q \in \mathbb{R}^3$ of each node x_i^q , $i=1..n_q$, in the fine grid \mathcal{M}^q . As n_q is potentially very large and the stiffness matrix \mathbf{A}^q very ill-conditioned due to the inhomogeneous medium, we wish to find a hierarchical approach to solving for the displacements which offers both better conditioning and homogenization properties, in the sense that coarse levels capture coarse approximations; consequently, this construction can be used for model reduction later on.

3.2 Matrix-valued canonical multiresolution analysis

While scalar basis functions are the default choice in typical finite element elasticity computations, we use 3×3 matrix-valued basis functions instead: since displacements are 3D vectors, coarse basis functions need to provide pointwise linear transformations so that the x , y , and z coordinates are not treated independently [Chen et al. 2018]. Anisotropy induced by the heterogeneity of the deformable body can then be properly captured even at coarse scales.

Finest functional space. We first derive nodal, matrix-valued basis functions $\{\varphi_i^q\}_i$ on the fine grid from the trilinear scalar functions:

$$\varphi_i^q = \mathbb{I}_{3 \times 3} \bar{\varphi}_i \quad (5)$$

where $\mathbb{I}_{p \times q}$ denotes the $p \times q$ diagonal matrix with unit diagonal values. These basis functions define a finite-dimensional functional space $\mathcal{V}^q = \text{span}\{\varphi_i^q\}_i$, used to reconstruct the continuous deformation in our framework from the node displacements u_i^q . Note that for this finest resolution, the associated stiffness matrix is the traditional trilinear stiffness matrix \mathbf{A}^q mentioned above.

Hierarchy of functional spaces. Our use of trilinear basis functions in Eq. (5) allows us to define a whole hierarchy of refinable (matrix-valued) basis functions on the sequence of meshes \mathcal{M}^k . Indeed, trilinear basis functions are known to be refinable on the mesh hierarchy we defined in Sec. 3.1: trilinear basis functions on level $q-1$ are, in fact, a linear combination of the finer trilinear basis functions $\bar{\varphi}_i$ —and similarly for the levels above. Therefore, we can build upon our matrix-valued basis functions φ_i^q to induce a hierarchy of basis functions φ_i^k for $k=1..q-1$ through the refinability relation:

$$\forall k \in \{1, \dots, q-1\}, \quad \varphi_i^k = \sum_{j=1}^{n_{k+1}} \mathbf{C}_{ij}^k \varphi_j^{k+1}, \quad (6)$$

where \mathbf{C}^k is a $3n_k \times 3n_{k+1}$ matrix for which $\mathbf{C}_{ij}^k = c_{ij} \mathbb{I}_{3 \times 3}$, and the scalar coefficients c_{ij} indicate the refinement relationship for the original trilinear basis functions (in 1D, these coefficients are $\frac{1}{2}, 1, \frac{1}{2}$, see inset). Sparsity of \mathbf{C}^k follows from the actual locality of the refinement stencil of trilinear functions. Note that this definition is a seemingly complicated way to explain that the resulting basis functions at coarser levels look exactly like the the basis functions at the finest level, just on coarser grids; but we will leverage the refinement matrices \mathbf{C}^k (and even propose simpler ones) to compute our material-adapted multiresolution analysis, hence the necessary introduction of this overly-pedantic notation. The resulting basis functions form a whole hierarchy of functional spaces \mathcal{V}^k , $1 \leq k < q$. As discussed in Sec. 2, these canonical basis functions lack eigen-awareness at coarse scales: applying them *as is* for Galerkin projection can be arbitrarily bad.

Kernel matrix. For each refinement matrix \mathbf{C}^k , $k=1..q-1$, we also define an associated sparse $3\eta_k \times 3n_{k+1}$ kernel matrix \mathbf{W}^k satisfying

$$\mathbf{C}^k \mathbf{W}^{k,T} = \mathbf{0}_{3n_k \times 3\eta_k}.$$

By its very definition, the rows of \mathbf{W}^k parametrize the kernel of \mathbf{C}^k ; therefore, each kernel matrix at level k provides a basis for elements of \mathcal{V}^{k+1} that *cannot be captured* at all in (i.e., that will be invisible to) \mathcal{V}^k . As such, these matrices will be crucial to the creation of material-adapted wavelets.

3.3 Towards material-adapted basis functions

We must now *adapt* the basis functions of the canonical hierarchy of functional spaces to form, based on the actual material of the deformable body \mathcal{B} , a material-adapted multiresolution analysis.

Refinability. One of the key arguments we made in Sec. 2 is that having refinability allows for adaptive computations. As a consequence, we propose to construct material-adapted refinement operators \mathbf{C}^k for each level k between 1 and $q-1$, from which the

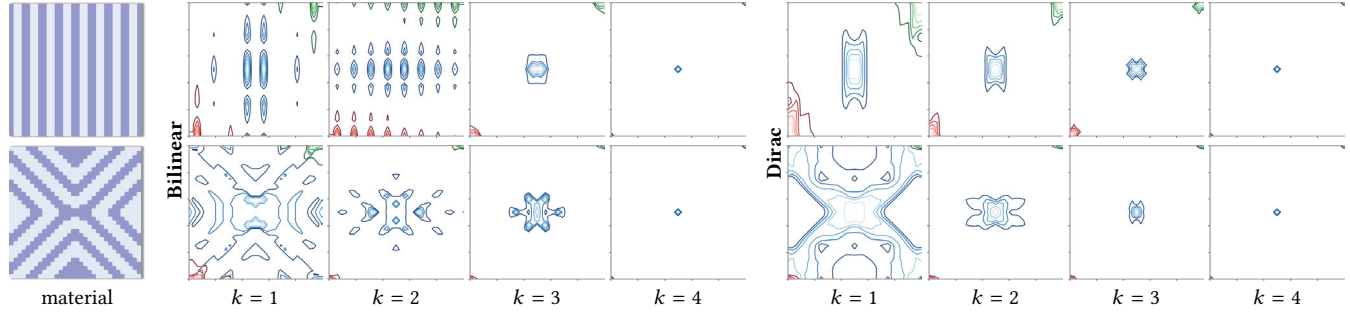


Fig. 4. **Material-adapted basis functions.** We visualize our matrix-valued basis functions ϕ^k (using isolevels of their *Frobenium norms*) for two different heterogeneous 2D materials (top/bottom) (no boundary conditions are set). The basis functions of three different node locations (center (blue), and two corners (red and green)) are shown on four different levels. At the finest level $k=q=4$, basis functions are bilinear. Coarser levels depend on whether Dirac (left) or bilinear (right) refinement is used, but the anisotropy of the material is clearly reflected in all the resulting basis functions.

material-adapted basis functions are simply defined through:

$$\phi_i^k = \sum_{j=1}^{n_{k+1}} C_{ij}^k \phi_j^{k+1}, \quad (7)$$

with $\phi_i^q \equiv \varphi_i^q$ for $i = 1..n_q$ to bootstrap the construction. In other words, we keep the finest functions unchanged, but the coarser functions are assembled from linear combinations of these fine basis functions. Compared to the original basis functions, they will remain sparse (with a local support decreasing in size with k), but will now span material-adapted functional spaces \mathcal{V}^k such that fundamental frequencies of scale k are well captured.

Wavelets. As a byproduct of the construction, we will also derive wavelets¹ (denoted as $\{\psi_j^k\}_j$ at level k , with $1 \leq k < q$) that are \mathcal{L} -orthogonal to the material-adapted basis functions, i.e.,

$$\int_{\Omega} \phi_i^k \mathcal{L} \psi_j^{k,T} = 0 \quad \forall i, j, \quad (8)$$

as in the scalar-valued case of Eq. (1).

3.4 Variational definition of material-adapted subsampling

We leverage the variational formulation of [Owhadi 2017] given in Eq. (2) by adapting it to our matrix-based functional spaces. Denoting \mathbb{A}^{k+1} the yet-unknown stiffness matrix computed with material-adapted basis functions of level k assembled from 3×3 blocks

$$\mathbb{A}_{ij}^{k+1} = \int_{\Omega} \phi_i^{k+1} \mathcal{L} \phi_j^{k+1,T},$$

the material-adapted \mathbb{C}^k can be defined variationally through

$$\mathbb{C}^k = \arg \min_M \text{Tr} [M \mathbb{A}^{k+1} M^T] \quad \text{s. t.} \quad M \mathbb{C}^{k,T} = \mathbb{I}_{3n_k \times 3n_k}. \quad (9)$$

Mirroring the scalar case given in its continuous form in Eq. (2), this definition looks for a new refinement operator at level k such that the material-adapted basis functions it induces minimize the quadratic form corresponding to the operator \mathcal{L} (given by its stiffness matrix \mathbb{A}^{k+1} computed at level $k+1$), under the constraint that they are weakly collocated with the canonical basis functions ϕ_i^k of the same level (i.e. they should be similarly positioned in physical space).

¹They are technically prewavelets since we will not orthonormalize them; but we will refer to them as wavelets for simplicity.

This constrained quadratic optimization problem turns out to have a closed form solution, reading:

$$\mathbb{C}^k = \mathbb{C}^{k,\dagger} \left[\mathbb{I}_{3n_{k+1} \times 3n_{k+1}} - \mathbb{A}^{k+1} \mathbf{W}^{k,T} \left(\mathbb{B}^k \right)^{-1} \mathbf{W}^k \right], \quad (10)$$

where $\mathbb{C}^{k,\dagger} = (\mathbb{C}^k \mathbb{C}^{k,T})^{-1} \mathbb{C}^k$ (i.e., the Moore-Penrose (left) pseudoinverse of \mathbb{C}^k) and $\mathbb{B}^k = \mathbf{W}^k \mathbb{A}^{k+1} \mathbf{W}^{k,T}$.

PROOF. Given the constraint $\mathbb{C}^k \mathbb{C}^{k,T} = \mathbb{I}_{3n_k \times 3n_k}$ and the fact that the columns of the kernel matrix \mathbf{W}^k span the kernel of \mathbb{C}^k , the solution refinement matrix \mathbb{C}^k must be of the form

$$\mathbb{C}^k = \mathbb{C}^{k,\dagger} + \mathbb{Z} \mathbf{W}^k \quad (11)$$

where \mathbb{Z} is an arbitrary $3n_k \times 3n_k$ matrix; that is, it must be equal to the pseudoinverse of \mathbb{C}^k plus a component lying in the kernel of \mathbb{C}^k . Moreover, to minimize the quadratic form \mathbb{A}^{k+1} while satisfying the constraint, their respective derivatives must be proportional, i.e., $\mathbb{C}^k \mathbb{A}^{k+1} = \lambda \mathbb{C}^k$ where λ is a Lagrange multiplier. In other words, $\mathbb{C}^k \mathbb{A}^{k+1}$ must be orthogonal to the kernel operator \mathbf{W}^k :

$$\mathbb{C}^k \mathbb{A}^{k+1} \mathbf{W}^{k,T} = \mathbf{0}. \quad (12)$$

Therefore, $(\mathbb{C}^{k,\dagger} + \mathbb{Z} \mathbf{W}^k) \mathbb{A}^{k+1} \mathbf{W}^{k,T} = \mathbf{0}$, which implies:

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

Substituting \mathbb{B}^k for the matrix product in parenthesis and plugging \mathbb{Z} back into the expression of \mathbb{C}^k leads to Eq. (10). \square

This explicit expression for \mathbb{C}^k has important consequences: it implies a simple and efficient fine-to-coarse construction of material-adapted basis functions and stiffness matrices entirely through sparse linear algebra as we will soon present.

Operator-adapted wavelets. Operator-adapted wavelets $\{\psi_i^k\}_{i,k}$ associated with the operator-adapted basis functions can also be directly derived through the kernel matrices via:

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

Indeed, these wavelets are \mathcal{L} -orthogonal to all basis functions on the same level k since, by definition, they satisfy:

$$\int_{\Omega} \phi_i^k \mathcal{L} \psi_j^{k,T} = \left(\mathbb{C}^k \mathbb{A}^{k+1} \mathbf{W}^{k,T} \right)_{ij},$$

which we proved to be a null matrix in Eq. (12) for any $i = 1..n_k$ and $j = 1..n_{k+1}$. We will denote the space of material-adapted wavelets

at level k by $\mathcal{W}^k = \text{span}\{\psi_i^k\}$ for $i=1..\eta_{k+1}$ in the remainder of this paper. Note also that the matrix \mathbb{B}^k used in Eq. (10) can now be understood as the Galerkin stiffness matrix of \mathcal{L} when written in the basis of wavelets, i.e., $\mathbb{B}_{ij}^k = \int_{\Omega} \psi_i^k \mathcal{L} \psi_j^{k,T}$.

Material-adapted multiresolution analysis. We now have a hierarchy of basis functions and wavelets from level 1 to level $q-1$, with a cardinality equal to the number of basis functions at the finest level q , offering a material-adapted multiscale framework for numerical simulation. In particular, while the discrete operator representing the elasticity operator \mathcal{L} using the fine functional space \mathcal{V}^q was the stiffness matrix \mathbb{A}^q , using our material-adapted basis instead (Eq. (4)) leads to a *block-diagonal* stiffness matrix \mathbb{L} :

$$\mathbb{L} = \begin{pmatrix} \mathbb{A}^1 & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbb{B}^1 & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbb{B}^{q-1} \end{pmatrix},$$

due to the orthogonality of the basis functions and wavelets. In addition, the properties of the variational definition of operator-adapted wavelets from [Owhadi 2017] imply that the diagonal blocks of this matrix are guaranteed to be well-conditioned, offering a better (still sparse) substitute to the fine-resolution matrix \mathbb{A}^q . We refer the reader to the recent work of [Budninskiy et al. 2019] for a short summary of the proof of this statement for the scalar-valued case, which still holds in our vector-valued case.

3.5 Algorithm

We now explain how having a closed form solution to our variational definition of material-adapted basis functions allows us to construct the entire hierarchy of basis functions, wavelets, and stiffness matrices in a fine to coarse fashion.

Storage of adapted basis functions and wavelets. It is clear from Eqs. (7) and (8) that all adapted functions are *linear combinations* of the finest basis functions φ_i^q . Therefore, for notational convenience, we can describe them as matrix coefficients of the fine basis functions they are made of: we use a (sparse) matrix Φ^k of operator-adapted basis functions (with $3n_k$ rows and $3n_q$ columns) and a (sparse) matrix Ψ^k of wavelets (with $3\eta_k$ rows and $3n_q$ columns) with

$$\Phi^k = \begin{bmatrix} \varphi_1^k \\ \varphi_2^k \\ \vdots \\ \varphi_{n_k}^k \end{bmatrix}, \quad \Psi^k = \begin{bmatrix} \psi_1^k \\ \psi_2^k \\ \vdots \\ \psi_{\eta_k}^k \end{bmatrix}$$

where each entry φ_i^k or ψ_j^k is discretized as 3×3 matrix coefficients to apply to the n_q basis functions of the fine mesh.


Fine-to-coarse computations at a glance. From the finest basis functions spanning \mathcal{V}^q as well as the refinement matrices \mathbb{C}^k (along with their precomputed pseudoinverse) and the kernel matrices, we can now assemble all relevant matrices of level $q-1$ in the order indicated in Alg. 1; the same order can then be used to deal with level $q-2$, etc, all the way down to the coarsest level 1. However, this approach is not efficient *as is*: as we mentioned earlier, the variational definition of the adapted basis functions $\{\varphi_i^k\}_i$ make

them decay exponentially fast, but it does not mean that they are actually sparse. We, in fact, need to sparsify both \mathbb{C}^k and \mathbb{A}^k ; we use simple truncation, setting to zero values below 10^{-9} .

Fast computation of \mathbb{C}^k . One key part in the algorithm is the evaluation of the intermediate matrix \mathbb{Z} (see Eq. (10) used in proof above) in line 3. Indeed, we compute a material-adapted refinement matrix in two stages: first, we perform a linear solve to obtain \mathbb{Z} , and then invoke Eq. (11) to find \mathbb{C}^k . Because $\mathbb{C}^{k,\dagger}$ can be precomputed offline, computational complexity is dominated by the first step, and the sparsity of \mathbb{Z} is crucial to the efficiency of the rest of the computations. Extending the approach of [Owhadi 2017] for the scalar-valued case, we find \mathbb{Z} verifying

$$\mathbb{B}^k \mathbb{Z}^T = -\mathbb{W}^k \mathbb{A}^{k+1} \mathbb{C}^{k,\dagger T}$$

by reducing (i.e., localizing) this linear system as follows. Since every column \mathbb{Z}_i of \mathbb{Z} corresponds to the i -th node of the coarser mesh \mathcal{M}^{k-1} , consider a small region R for which all nodes are less than ρ nodes away from this i -th node on the computational grid \mathcal{M}^{k-1} . Let S be the set of indices such that $\ell \in S$ if and only if the ℓ -th row of \mathbb{W}^k has at least one non-zero entry on nodes of the finer mesh \mathcal{M}^k enclosed in R . Consider a submatrix \mathbb{b} of the stiffness \mathbb{B}^k that includes only the rows and columns whose indices are in S ; this is the “reduced” stiffness matrix of the wavelets around node i whose local support intersects R . Consider also a vector \mathbf{z}_i that subsamples the i -th column of $\mathbb{W}^k \mathbb{A}^{k+1} \mathbb{C}^{k,\dagger T}$ to include only the indices from S . The entries of column \mathbb{Z}_i with indices from S , assembled in a vector \mathbf{z}_i , are computed by solving a small linear system $\mathbb{b} \mathbf{z}_i = -\mathbf{z}_i$, while remaining entries of \mathbb{Z}_i are set to zero. In practice this approach to computing \mathbb{Z} is considerably faster than the full linear solve for large matrices, and is particularly easy to parallelize. By construction it also promotes sparsity for \mathbb{Z} , and in turn, for the adapted refinement matrix \mathbb{C}^k and basis functions. In practice, we pick $\rho=3$, which corresponds to performing local solves only in 3-ring neighborhoods around each element; we will discuss other possible tradeoff in Sec. 3.7.



(a) groundtruth (b) with corrections (c) without corrections
 Fig. 5. **Importance of Geometric Invariance.** Even for a moderately large sparsification radius $\rho=5$, the induced loss of infinitesimal translation and rotation invariance can have drastic homogenization consequences (right), unless our simple local correction is applied (middle).

Corrections to enforce invariance. Using localized solves to compute \mathbb{Z} promotes sparsity in \mathbb{C}^k . However, this forced sparsity also breaks the infinitesimal invariance to translation and rotation of the adapted basis functions (see Fig. 5, where sparsification during the solve of a simple shear or a simple twist grossly deforms the result, unless correction is applied)—an issue that failed to be identified in [Owhadi 2017; Owhadi and Zhang 2017] where a similar treatment loses partition of unity in the scalar case. In order to ensure that constant displacements and infinitesimal rotations are preserved after sparsification, one should impose:

$$\forall j, \sum_i \mathbb{C}_{ij}^k = \mathbb{I}_{3 \times 3}, \quad \forall j, \sum_i \mathbb{C}_{ij}^k [\bar{\mathbf{x}}_i^{k-1}]_{\times} = [\bar{\mathbf{x}}_j^k]_{\times},$$

where $[\cdot]_{\times}$ denotes the 3×3 skew-symmetric matrix encoding cross product, and $\bar{\mathbf{x}}_i^k$ denotes the position of node i at level k of the object \mathcal{B} at rest; one recognizes the “geometric” conditions of adapted basis functions used in [Chen et al. 2018], written this time using our notation. Without these conditions, spurious aphysical modes pollute the solution space. We enforce these constraints in a post-processing stage right after \mathbb{C}^k has been evaluated at level k by finding, for each node j at level k , the smallest (in Frobenius norm) changes to the submatrix of \mathbb{C}^k associated with that node which enforce these two conditions and keep the sparsity pattern intact. That is, for each node j at level k and calling \mathcal{I} the set of indices i of mesh nodes at level $k-1$ with non-zero submatrix \mathbb{C}_{ij}^k , we compute 3×3 adjustment matrices $\{\delta_i\}_{i \in \mathcal{I}}$ as the minimizers of $\sum_{i \in \mathcal{I}} |\delta_i|_F^2$ subject to the two constraints $\sum_{i \in \mathcal{I}} \delta_i = \epsilon_1$ and $\sum_{i \in \mathcal{I}} \delta_i [\bar{\mathbf{x}}_i^{k-1}]_{\times} = \epsilon_2$, where ϵ_1 and ϵ_2 are the 3×3 matrix residuals of the two geometric conditions above for node j . (We use the Schur complement to efficiently compute the solutions of this small constrained minimization.) If none of the indices in \mathcal{I} correspond to a node with Dirichlet boundary conditions, we then set $\mathbb{C}_{ij}^k \leftarrow \mathbb{C}_{ij}^k + \delta_i$ for all $i \in \mathcal{I}$ since these invariances have to hold only *away* from fixed regions.

<p>Input: Basis matrix Φ^q and stiffness matrix $\mathbb{A}^q \equiv \mathbb{A}^q$ on finest level, refinement/kernel matrices $\{\mathbb{C}^k, \mathbb{W}^k\}_{k=1}^{q-1}$.</p> <p>1 for $k \leftarrow q-1$ to 1 do</p> <p>2 Compute wavelet stiffness: $\mathbb{B}^k \leftarrow \mathbb{W}^k \mathbb{A}^{k+1} \mathbb{W}^{k,T}$;</p> <p>3 Compute matrix \mathbb{Z} with fast localized solves via $\mathbb{B}^k \mathbb{Z}^T = -\mathbb{W}^k \mathbb{A}^{k+1} \mathbb{C}^{k,\dagger T}$;</p> <p>4 Compute adapted basis refinement matrix: $\mathbb{C}^k \leftarrow \mathbb{C}^{k,\dagger} + \mathbb{Z} \mathbb{W}^k$;</p> <p>5 Compute adapted wavelets on level $(k-1)$: $\Psi^k \leftarrow \mathbb{W}^k \Phi^{k+1}$;</p> <p>6 Compute adapted basis on level k: $\Phi^k \leftarrow \mathbb{C}^k \Phi^{k+1}$;</p> <p>7 Assemble stiffness matrix for lower resolution level: $\mathbb{A}^k \leftarrow \mathbb{C}^k \mathbb{A}^{k+1} \mathbb{C}^{k,T}$ (sparsified through truncation);</p> <p>8 end</p> <p>Output: Material-adapted bases $\{\Phi^k\}_k$, wavelets $\{\Psi^k\}_k$, stiffness $\{\mathbb{A}^k\}_k$, $\{\mathbb{B}^k\}_k$ and refinements $\{\mathbb{C}^k\}_k$.</p>

Algorithm 1: Adapted Basis and Wavelet Construction

3.6 Solution from adapted hierarchy

With the hierarchical procedure described above, we now have material-adapted stiffness matrices \mathbb{A}^k and \mathbb{B}^k for $1 \leq k < q$. Instead of solving for the fine displacements of the single-resolution elastostatics problem through $\mathbb{A}^q u^q = g^q$ (where g^q encodes forcing), our multiscale approach allows us to use material-adapted variables

instead as the fine functional space has been decomposed into a whole hierarchy (see Eq. (4)): one must solve the following q smaller sparse linear systems:

$$\mathbb{B}^k \mathbb{w}^k = \mathbb{W}^k \mathbb{g}^{k+1} \text{ for } q-1 \geq k \geq 1 \quad \text{and} \quad \mathbb{A}^1 \mathbb{v}^1 = \mathbb{g}^1,$$

where g^q is coarse-grained through $\mathbb{g}^k = \mathbb{C}^k \mathbb{g}^{k+1}$ with $\mathbb{g}^q \equiv g^q$, the vectors \mathbb{w}^k store the wavelet coefficients between levels. The fine solution u^q is then recovered through:

$$u^q = \Phi^{1,T} \mathbb{v}^1 + \sum_{k=1}^{q-1} \Psi^{k,T} \mathbb{w}^k,$$

reflecting the fact that our basis functions and wavelets block-diagonalize the elasticity operator. One can incorporate only a subset of the wavelets too if details are not needed.

3.7 Discussion

The construction we just presented is quite general for vector-valued PDEs, and has a few parts that one can choose differently based on the type of application sought after.

Choice of canonical refinement. While we used the natural refinement \mathbb{C}^k of trilinear basis functions, one can in fact use any refinement relation for known refinable functions, even if we keep trilinear basis functions at the finest level: consistency between the two is not formally needed. In the context of computer graphics where efficiency is paramount, one may in fact prefer the *sparsest* possible refinement, where a coarse basis function ϕ_i^k corresponds to a single basis functions at level $k+1$ at the same location (a concept sometimes referred to as *lazy wavelet*). This “Dirac” refinement, recently proposed in [Budninskiy et al. 2019], further simplifies the construction since the refinement matrix becomes sparsest, and its kernel is particularly trivial too; however, it comes at a price: the approximation properties are weakened (Fig. 9), even if the condition numbers of the resulting stiffness matrices are actually improved. That is, coarse-graining with the Dirac refinement is less accurate for the same number of levels. However, we found that the Dirac refinement is often sufficient for animation purposes given the appreciable amount of computational time it saves. It has the added benefit that coarse displacements are precisely subsampling fine displacements, an advantageous property in various applications.

Choice of sparsification. Our approach to the sparsification of \mathbb{Z} employed a support restricted to a radius of ρ elements around each coarse node. The choice of an integer ρ is, of course, non-unique (it may even be changed on a per-level basis), but its role is clear: taking too large a ρ will lead to less sparse stiffness matrices \mathbb{A}^k and \mathbb{B}^k , while forcing ρ to be too small may lose significant accuracy. This is the traditional *trade-off* between spatial vs. eigenspatial localization, akin to the Heisenberg uncertainty principle. In our experience, picking $\rho = 3$ is enough for visual purposes, but setting ρ to 4 or 5 for heterogeneous materials with strong contrast between their constitutive materials clearly improves accuracy. If accuracy is paramount, then sparsification through thresholding by picking a cut-off value based on a chosen percentile is, of course, preferable. Fig. 6 shows how the homogenization error depends smoothly on both the sparsification radius and the cut-off value.

4 MODEL-REDUCED LINEAR ELASTICITY

We now discuss how to apply the material-adapted approach for the efficient simulation of linear elastic inhomogeneous materials.

4.1 Elastostatics

Once an object \mathcal{B} has been discretized on a fine grid \mathcal{M}^q , we can directly compute its stiffness matrix \mathbb{A}^q using trilinear basis functions accounting for the inhomogeneous material. The coarse-graining procedure described above can then be used to compute the material-adapted stiffness matrices \mathbb{A}^k and \mathbb{B}^k for $1 \leq k < q$, providing an alternative to solving the usual fine-resolution elastostatics problem written in \mathcal{V}^k via $\mathbb{A}^q u^q = g^q$ as described in Sec. 3.1.

More interestingly, we can now approximate the solution based on any coarse level k of our hierarchy: if we wish to limit computations to a given level k , we can simply solve a single sparse linear system

$$\mathbb{A}^k \mathbf{u}^k = \mathbf{g}^k,$$

(where \mathbf{g}^k is directly computed from g^q as described in Sec. 3.6) or, if we pre-factorized \mathbb{A}^k , through fast forward- and back-substitutions. Because the stiffness matrix at a given level is designed to properly capture the fundamental frequencies at this scale, the coarse displacements \mathbf{u}^k on \mathcal{M}^k will be predictive of the general solution. We can display a fine solution from these coarse displacements: $u^q = \Phi^{k,T} \mathbf{u}^k$ will be a good approximation of the full solution due to the coarse-graining nature of our construction. Note that one can also decide to use the first m levels instead, to improve accuracy without raising the complexity significantly.

4.2 Dynamics

Until now, we have mostly discussed the case where the operator \mathcal{L} is involving purely spatial derivatives. Yet, material-adapted basis functions can also deal seamlessly with dynamics by not only coarse-graining a spatial operator, but its integration in time as well. Indeed, a number of integrators for linear elasticity can be written as either $\dot{u}_{t+dt} = \mathcal{L} \dot{u}_t$ (thus performing explicit integration) or $\mathcal{L} \dot{u}_{t+dt} = \dot{u}_t$ (implicit integration). Now the linear operator \mathcal{L} can be seen as a discrete flow of the mechanical system. Modulo the fact that the operator is no longer just the elasticity operator but its dynamic flow in time, the exact same treatment of the full solution of the time update via our material-adapted wavelets or through a single coarse scale applies. We will show that the typical phase shift observed in previous works is greatly reduced, without having resort to solutions involving, e.g., the tuning of the actual material stiffness [Chen et al. 2017] that do not apply to inhomogeneous materials.

4.3 Non-linear geometric warping

Linear elasticity exhibits significant visual artifacts when large rotations are involved, reflecting the fact that this linearized variant of elasticity is based on a strain tensor that is *not* invariant under rotations. In finite-element based approaches to linear elasticity, the use of corotational (CR) methods has been proven an efficient way to virtually remove this issue at runtime by deriving a local rotation and expressing the local strain in the corotated configuration [Müller et al. 2002; Hauth and Strasser 2004; Kugelstadt et al. 2018]. For modal analysis, the linear map from modes to Euclidean




Fig. 6. **Error & density w.r.t. sparsification radius and cut-off.** For the example in Fig. 5, we plot the homogenization error $\log(\|u - u^1\|_{\mathcal{L}} / \|g\|_{L^2})$ and density (i.e., percentage of non-zero entries in \mathbb{C}) as a function of both sparsification radius ρ and (log of the) cut-off threshold.

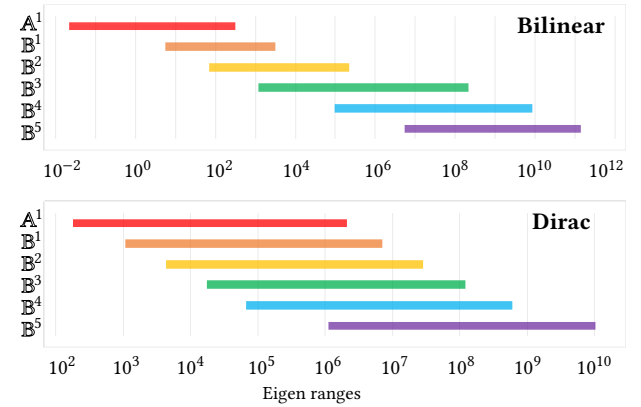


Fig. 7. **Eigenvalue ranges of stiffness matrices.** For the inhomogeneous material from Fig. 2, the range of generalized eigenvalues for the wavelet matrices and the coarsest stiffness matrix of our material-adapted hierarchy covers overlapping frequency bands of the original fine stiffness matrix.

coordinates is also notoriously inadequate for large deformation, so modal warping [Choi and Ko 2005] and rotation-strain (RS) coordinates [Huang et al. 2011] were proposed to reconstruct plausible shapes from modes.

Non-linear post-warping. Since we are able to simulate linear elasticity for complex materials on coarse grids, it is tempting to use a corotational simulation at runtime to palliate the visual effects of linear elasticity. However, it is no longer trivial to derive a good corotated frame field. Typical approaches rely either on the piecewise constant nature of the gradient of displacement to perform polar decomposition [Hauth and Strasser 2004] or a per-element minimization of the displacement magnitudes [Georgii and Westermann 2008]; neither options are particularly simple to extend reliably to our inhomogeneous case with adapted basis functions. Instead, we found RS coordinates, typically used in modal analysis, provide us with a simpler alternative which already provide better results than existing coarse-graining approach. RS coordinates transform a local gradient of displacement G by exponentiating the antisymmetric part of the gradient (to extrapolate a large rotation) while the strain is restricted to the symmetric part of G plus identity. For computational efficiency, we instead use:

$$RS(G) = \text{Cay}[(G - G^T)/2] (\mathbb{I}_{3 \times 3} + (G + G^T)/2) - \mathbb{I}_{3 \times 3},$$




Fig. 8. **Conditioning.** For the same example as in Fig. 2, we display the condition numbers of the stiffness matrices computed with the material-adapted basis functions and wavelets in the hierarchy. The matrices A^k get better conditioned as one approaches the coarsest level 1, while the matrices B^k remain well conditioned throughout.

where the Cayley map $\text{Cay}(\cdot)$ applied to a skew-symmetric matrix A is defined as the rotation $\text{Cay}(A) = (\mathbb{I}_{3 \times 3} - A/2)^{-1}(\mathbb{I}_{3 \times 3} + A/2)$, a much more efficient alternative to the exponential map [Kobilarov et al. 2009]. To inject the desired non-linearity in the result of our simulations, we first express our coarse-grained solution u^k at a chosen level k as fine displacements u_i^q through $u^q = \Phi^{k,T} u^k$; we then find the “corotated” displacement \hat{u}^q by solving

$$\hat{u}^q = \arg \min_u \int_{\Omega} \|\nabla u - RS(\nabla u^q)\|_F^2 \quad \text{s. t.} \quad S\hat{u}^q = 0, \quad (13)$$

where S is the selection matrix that spatially fixes the fine nodes that are subject to Dirichlet boundary conditions during the animation. The use of RS-based transforms and Poisson reconstruction [Huang et al. 2011] creates a non-linear mapping which drastically removes the linearization artifacts of our model. As an added bonus, we noticed that the usual “over-estimation” of the rotation part present in RS coordinates is significantly reduced by replacing the exponential by the Cayley map, a simpler approach than the use of Padé approximants recommended in [Pan et al. 2015]. Finally, if only a *coarse* representation is desired, we can proceed even more efficiently: since the use of a Dirac refinement precisely subsamples the fine displacements, we can perform the Poisson reconstruction and RS transforms at *any* level k directly. This was done at the coarsest level in Fig. 12 for instance, demonstrating that the non-linearity injection procedure behaves properly at various scales.

Fast implementation. In our implementation, 8 Gaussian cubature points are used in each hexahedral element (and one cubature per tetrahedral element if a simplicial mesh is used), on which the deformation gradient and its RS transform are evaluated in parallel. The final sparse linear system solving Eq. (13) is pre-factorized for efficiency. If collisions happen during the simulation, the Poisson matrix can be updated to constrain new nodes via low-rank updates.

5 RESULTS AND COMPARISONS

We now present various results to demonstrate the value of our approach on heterogeneous materials.

5.1 Mesh hierarchy

While we used nested regular grids to explain our approach, its construction is very general and does not make any special assumption on how these meshes are related (for instance, refinements can be nested or not). Due to the large body of known subdivision schemes for both simplicial and polyhedral meshes of arbitrary domains, any hierarchy of meshes for which a node-based subdivision




Fig. 9. **Homogenization error plots.** For a composite object made out of juxtaposed layers of soft (Young’s modulus $E = 100$) and stiff ($E = 5 \times 10^6$) materials (using the two patterns (left/right) of Fig. 4), we plot the homogenization error $\|u^k - u\|_{\mathcal{L}} / \|g\|_{\mathcal{L}_2}$ for the five intermediary levels k (finest: 32×32 ; coarsest: one element) of our coarsening algorithm. Bilinear refinement is more accurate, but both slopes are similar.

scheme exists could be used in our construction. However, regular (Cartesian) grids are undeniably simpler due to their tensor product nature, and often more amenable to efficient implementation. We show examples on tetrahedral meshes in Figs. 1, 5, 13 and 18, while all other examples are performed on regular grids.

5.2 Visualizing material-adapted basis functions

Fig. 4 shows the effects of material inhomogeneity on the basis functions. While Dirac and trilinear refinements induce very isotropic elasticity-adapted basis functions (with, as expected, smaller support for the Dirac refinement case), a metamaterial made out of vertical strips of alternately stiff and soft materials has basis functions clearly exploiting the natural anisotropy created by this material sandwiching. The support remains local at all scales as expected. Fig 1(left) shows another 2D example of basis functions and wavelets visualization; notice that the Dirichlet boundary conditions imposed of the left side of the domain affects the shapes of the coarse basis functions as well (hence the non-symmetry of the coarsest basis function), as expected to offer coarse-graining.

5.3 Numerical conditioning and accuracy

To demonstrate the fact that each level of our hierarchy captures its own frequency “band” even for an inhomogeneous material, we display in Fig. 7 the ranges of non-zero eigenvalues for each stiffness matrix of a simple heterogeneous material. While the original matrix A^q on the fine scale covers a large spectrum due to the multiscale nature of its physical properties induced by material inhomogeneity, wavelet stiffness matrices B^k correspond to small ranges adapted to their level, while A^1 covers only the lowest range of the initial spectrum; these matrices are thus better conditioned.

In Fig. 8, we show the condition number of each stiffness matrix instead. Here again, we see a gain of several orders of magnitude for the hierarchical stiffness matrices (in particular, for the wavelet stiffness matrices) compared to the original fine one. As we go up in the hierarchy (from fine to coarse), the condition number of the stiffness matrices associated with coarser and coarser basis functions goes down as expected at each level.

Finally, we also demonstrate in Fig. 9 that the numerical error created by our coarse graining grows moderately as we go to coarser levels: the slopes we found for both trilinear and Dirac refinements follow the expected (optimal) error decay rate in “energy norm”

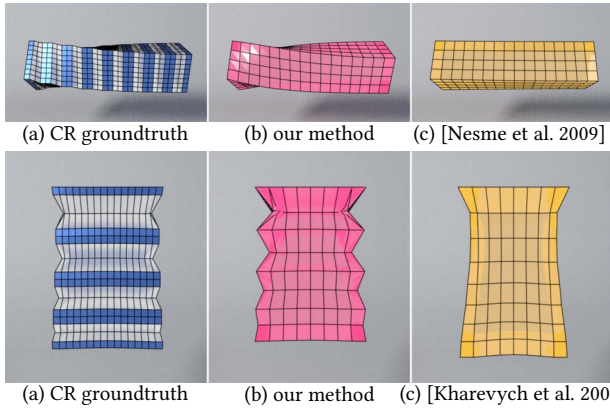


Fig. 10. **Comparisons.** Even for a single-level coarsening (from a 16^3 grid to a 8^3 grid), previous approaches can either fail to predict the correct stiffness (top) of a composite material, or only capture its rough behavior (bottom).

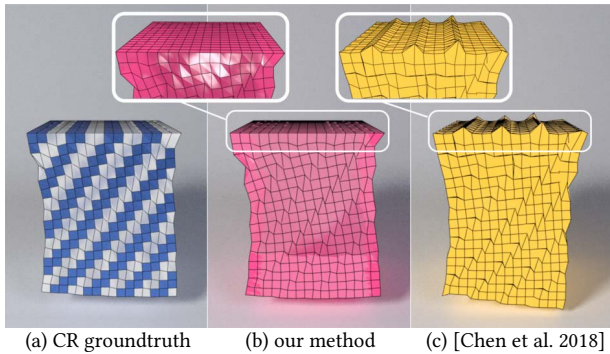


Fig. 11. **Boundary conditions.** When a composite material is hung from one of its faces, our coarse-graining naturally accounts for Dirichlet boundary conditions by leveraging the fine stiffness matrix. Instead, previous methods can exhibit large artifacts on (and away from) constrained parts.

(where the error is measured as $\int_{\Omega} (u - \hat{u})^T \mathcal{L}(u - \hat{u})$ and normalized by the L_2 norm of the forcing term [Owhadi 2017]). Moreover, we also see that the sparser Dirac refinement does indeed induce more errors, without changing the error slope. Note that none of the other homogenization methods [Nesme et al. 2009; Kharevych et al. 2009] fit on this graph, producing significantly larger errors as we will visually demonstrate next.

5.4 Coarse-graining compared to previous methods

Besides numerical evidence of the intrinsic value of our construction, we also present visual comparisons to previous work to show the comparative efficacy of our coarse graining procedure.

Comparisons to [Kharevych et al. 2009] and [Nesme et al. 2009]. The work of Kharevych et al. [2009] also offers coarse-graining of elasticity by homogenizing the material properties of each coarse element instead of designing special basis functions. While their approach is typically able to capture the coarse shape of a deformed body (for instance, for the hanging two-material example in Fig. 10(bottom)), it typically loses many of the detailed behavior compared to our approach. More similar to our method, Nesme et al. [2009]

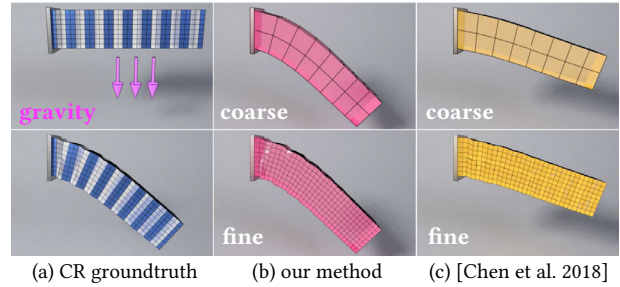


Fig. 12. **Spurious stiffness.** A pinned beam released under gravity, with two-level coarsening (64 times less elements). The approach in [Chen et al. 2018] ignores boundary conditions, leading to spurious structural rigidity (right) for such an aggressive coarsening. Our coarse simulation with Dirac coarse-graining (middle) captures nearly perfectly the fine simulation (top), and using our basis functions to reconstruct the fine deformed grid (bottom) shows the expected fine details of the original result (left).

proposed the design of basis functions to homogenize elastic properties. However, their construction makes various approximations such as junctions between elements being considered frozen; and indeed, we found that their results are systematically too stiff, see Fig. 10(top) for a clear example. Note that our coarse basis functions have a larger support, rendering our coarse stiffness matrix at the coarsest level less sparse than previous methods. However, since we prefactorize this matrix, the computational complexity at run time is not noticeably affected: we still gain two orders of magnitude speedup over computing directly the fine solution.

Comparisons to [Chen et al. 2018]. Our method is, in many ways, an extension of the work of Chen et al. [2018]: they also proposed a construction of matrix-based basis functions defined on the coarse grid (and stored on the fine grid) so that their span captures the behavior of an inhomogeneous material. We offer, however, several improvements over this work:

- first, our approach naturally captures boundary conditions at every level, whereas Chen et al. disregard boundary conditions to construct basis functions. This is achieved in our approach directly during the assembly of the fine stiffness matrix A^f , so the hierarchical construction carries this information throughout the levels. As Fig. 1 demonstrates, Dirichlet boundary conditions affect the actual shape of basis functions: the coarsest basis functions, for instance, clearly exhibit an asymmetry due to the left side being physically clamped. Failure to account for Dirichlet conditions does create strong artifacts, see Fig. 11 for an example. Moreover, as we further coarsen, ignoring boundary conditions creates spurious structural rigidity as demonstrated in Fig. 12.
- second, our principled construction extends their approach significantly. The quadratic form being minimized in our variational definition is directly related to the elasticity operator, and the imposed constraint enforces spatial collocation. Because we do not force the support to be only one coarse element wide, we also significantly improve the coarse graining accuracy. Moreover, we do away with the non-conforming nature of the basis functions, offering less artifacts in practice. Finally, because of the refinability we impose, our hierarchical construction involves

model (mesh type)	#nodes		#lvls	build	sparsity			storage	linear					CR		
	fine	coarse			\mathbb{B}^1	\mathbb{C}^1	\mathbb{A}^1		pref. A^q	solv. g^q	pref. A^1	solv. g^1	resid.	fine	coarse	speedup
letter H (tet)	15606	2656	2	11.5s	99.92%	98.8%	88.3%	300.3M	3.54s	0.08s	1.2s	0.011s	0.8	—	—	—
shark (hex)	18234	3317	2	15.1s	99.86%	97.79%	86.2%	390.5M	3.93s	0.10s	1.4s	0.012s	0.60	18.7s	0.036s	519.4
botijo (tet)	8026	1458	2	3.9s	99.85%	98.2%	86.7%	123M	0.95s	0.034s	0.2s	0.005s	0.11	10.7s	0.022s	486.36
hand (tet)	5672	929	2	2.4s	99.78%	97.1%	84.7%	83.5M	0.41s	0.023s	0.12s	0.002s	0.45	7.1s	0.011s	645.45
bird (tet)	5120	904	2	2.3s	99.76%	96.75%	79.8%	80.4M	0.62s	0.022s	0.14s	0.002s	0.55	6.3s	0.01s	630.0
square (quad)	4096	4	7	11.1s	0.0%	0.0%	0.0%	437.8M	0.079s	0.005s	6e-5s	4e-6s	0.09	—	—	—
beam (hex)	2673	81	3	15.3s	10.7%	24.4%	2.2%	320.8M	0.87s	0.008s	3e-3s	2.3e-4s	0.24	6.4s	0.003s	2133.3
dragon_1 (tet)	23811	3916	2	23.7s	99.95%	99.1%	92.98%	456.9M	8.0s	0.13s	1.7s	0.018s	1.1	—	—	—
dragon_2 (tet)	163152	3916	2	914s	99.992%	99.06%	93.28%	2.4G	336.6s	0.978s	2.04s	0.016s	2.4	—	—	—

Table 1. **Quantitative evaluation.** In this table, we provide quantitative evaluation of our method for different models undergoing static deformation. We include the size of the models, the time it took to construct the hierarchy (build), the sparsity of the resulting matrices (\mathbb{B}^1 , \mathbb{C}^1 and \mathbb{A}^1), and memory usage (to store matrices \mathbb{A}^k , \mathbb{B}^k , \mathbb{C}^k , \mathbb{W}^k , $\mathbb{C}^{\dagger,k}$, Φ^k and Ψ^k). Timings for the linear case are given for both the prefactorization part (pref.) and the back-substitution part (solv.), and the residual of the coarsest solve is also provided (resid.), evaluated as $\|u^1 - u\|_{\mathcal{L}} / \|g\|_{L_2}$. Timings for all our examples are indicated in seconds per frame. The sparsification radius here is set to be 5 and the cut-off threshold is 10^{-7} ; CHOLMOD [Davis 2009] was used to solve linear systems.

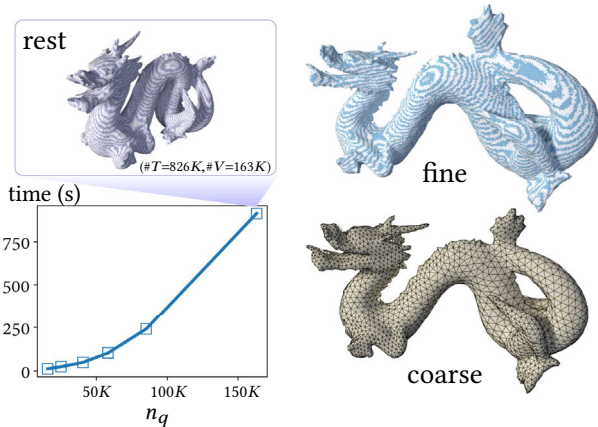


Fig. 13. **Constructing the material-adapted hierarchy.** For the dragon model, we plot the time it takes to construct one level of the material-adapted hierarchy (basis functions and wavelets) for various mesh sizes, from 2K to 160K vertices. The observed complexity matches the expected $O(n_q \log^7 n_q)$ of the scalar homogenization case from [Owhadi 2017].

sparse linear algebra computations for which a wide variety of numerical tools can be used; in their approach, a large number of local optimizations with a dedicated solver were required instead; note that these local optimizations grow quickly in size with the required amount of spatial coarsening. We can thus coarsen a given fine grid much more aggressively than they can for a given amount of computing power.

- third, we demonstrate preliminary results for the coarsening of dynamics; Fig. 18 shows that applying our coarse-graining procedure on an implicit Euler integration of linear elasticity shows a small phase shift (due to the forced sparsification of our procedure) compared to a fine simulation, reducing the temporal mismatch significantly compared to their quasi-static approach. Consequently, our approach is both more flexible and more efficient, and leads to visually more faithful coarse simulations.

5.5 Timing

Since we can coarsen a given fine mesh all the way down to a single element in our approach (see Fig. 16 and sixth row in Table 1; the 0% sparsity is due to the fact that it is coarsened to one element, hence resulting in a very small, but dense system), the computational gains at runtime compared to a fine simulation can be made




Fig. 14. **Effects of sparsification radius.** The material-adapted basis and wavelet construction (on letter H example) can get gradually slower if we decrease sparsity (i.e., if the sparsification radius grows), since the localization of computations worsens.

arbitrarily large. This may come, obviously, at the cost of accuracy; in particular, the use of Dirac refinement can exhibit artifacts if too coarse resolutions (and none of the wavelets) are used: see Fig. 2 where a 2×2 grid is used. The use of higher-order refinements would be even better at homogenizing the fine details, but at the price of a slower hierarchical construction due to the less sparse refinement relation across levels. For both Dirac and trilinear refinements our tests find, in agreement with [Chen et al. 2018], that three or more orders of magnitude speedups can be achieved with only limited visual artifacts—often with significant improvement compared to previous methods, see Figs. 10, 11 & 12 and Table 1 summarizing the various timings of our approach. Note that with our current unoptimized implementation, the entire five-level construction over a small 16×16 grid takes less than a second, while the large shark example (18K vertices, 10K elements) in Fig. 15 took less than 1 mins to precompute for a two-level coarsening on a plain Intel™ Xeon desktop. We notice that if we increase the sparsification radius ρ , precomputation time will rapidly increase (see Fig. 14), especially for hexahedral meshes — although the effect on the quality of coarse-graining for this example is barely noticeable. While we believe that the material contrasts we used are already exceeding the typical needs in computer animation, computational applications with stronger contrasts and where accuracy is key (thus needing a radius $\rho > 3$) would require optimizing the code for better memory management and for a more efficient use of parallelization. Finally, we point out that our CR runtime computations could be made more efficient by following the recent work of [Kugelstadt et al. 2018].

Observe also that like most coarsening (reduction) methods trying to build spectral approximations [Chen et al. 2018; Liu et al.




Fig. 15. **Jumping shark.** On this shark model, a coarse simulation of its body bending closely matches the fine version despite being eight times coarser, with an effective computational speedup factor of 519.

2019], our construction of material-adapted basis functions is computationally intensive: as Table 1 indicates, building the hierarchy takes from a few seconds to a few minutes. However, this offline process results in a well-behaved reduced system which, compared to solving a fine, badly-conditioned linear equation repeatedly, can accelerate simulation by potentially orders of magnitude, see Table 1. Our construction does not only allow for a faster solve, but it also provides material-adapted basis functions such that any external force field can be properly distributed to the coarse nodes for proper homogenization. Both features may thus be useful for a number of applications in animation, to efficiently compute reduced projective dynamics [Brandt et al. 2018] over a long sequence, or, as a multigrid preconditioner, to boost the efficiency of geometrically nonlinear elasticity simulation. Our method may also help to extract representative features of the operator to ease the training of large scale machine learning tasks, e.g., mechanical classification of tiling structures by neural network [Liu et al. 2019].

5.6 Mechanical analysis and structure analysis

In various applications such as inverse design [Schumacher et al. 2015] and mechanical characterization [Schumacher et al. 2018], computing the stress induced by given loads provides crucial information on the mechanical properties of the object’s geometric/material structure. Accuracy of the stress field has thus a significant impact on the evaluation and validation of feasibility for structural design, and fast evaluation is highly desirable. Fig. 16 demonstrates that our approach captures stress in complex composite materials with high accuracy even after very aggressive coarsening. Similarly, Fig. 17 illustrates the anisotropy of the resulting stiffness introduced by different tilings in a unit square domain: inspired by tests presented in [Schumacher et al. 2018] (note that they used periodic tilings, while we use a single square domain to account for boundary effects), we plot the directional Young’s modulus for four different composite materials. As expected, the physical behavior of the square changes based on the material distribution, inducing various stiffness changes along the diagonals and sides of the square exemplar.

5.7 Limitations

Coarse graining is still a recent topic, and the current interest in our animation field is bound to bring further improvements. We point out here what we see as the most obvious current limitations.




Fig. 16. **Stress visualization.** For two different bimaterial composite objects in 2D (with a contrast of 10^6 between soft (white) and stiff (black) material, left), a deformation is generated by fixing the top boundary of the domain and setting a downward surface traction on the bottom. Computing the resulting Piola-Kirchhoff stress tensor on all the cubature points for a fine simulation (on a 64×64 grid, middle) is nearly identical to the homogenized solution using one *single* element (right). Even for this aggressive coarsening, the stress distribution is homogenized very accurately.





Fig. 17. **Directional Young’s modulus.** The directional Young’s modulus is calculated as $E(d) = 1/\text{div}(dd^T) : (\mathbf{A}^1)^{\dagger} : \text{div}(dd^T)$, where \mathbf{A}^1 is the coarsest stiffness matrix and d is the direction (unit vector).

First, our postwarping approach is an efficient choice to deal with geometric non-linearity, but is not ideal if accurate large deformation are desirable: we may not match the local rotations of a fine corotational method. The inset shows a composite hand under gravity; while coarse and fine simulation of linear elasticity without corotational treatment match perfectly (yellow color, with the typical overstretching), our RS-corrected eight-times coarsened version fails to match the fine CR simulation — but still outperforms previous work. Second, our approach to sparsification is also perfectible: enforcing sparsity first, followed by a correction for geometric invariance, may not be optimal for efficiency. Finally, one may be able to find an approximation of our construction lending itself well to more efficient implementations: the hierarchical construction could then be performed at each (or every few) time step(s) to homogenize even non-linear models via linearization around the current shape.



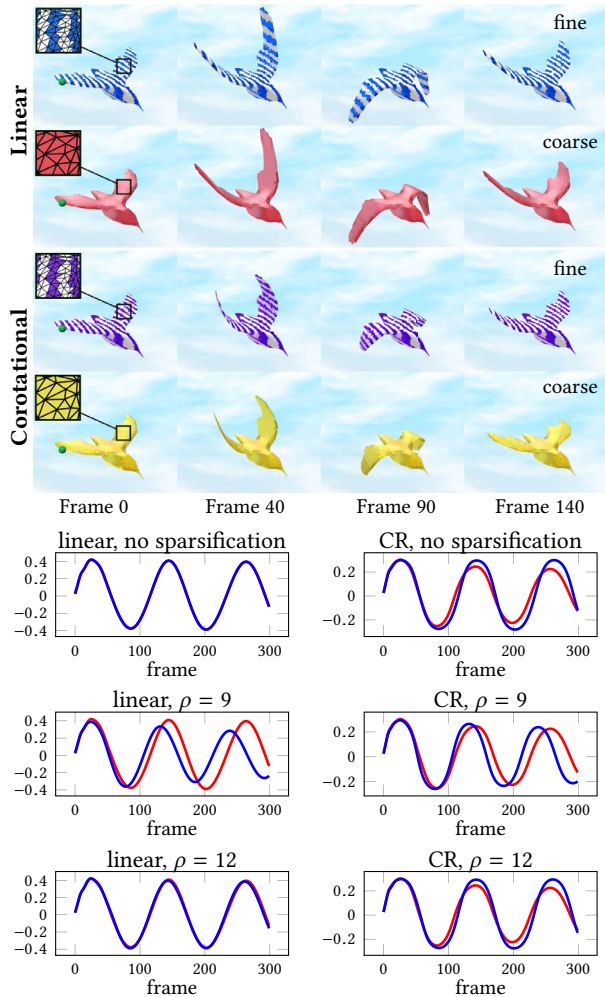


Fig. 18. **Coarse-graining of time integration.** We can apply our coarse-graining procedure to an integrator scheme (here, implicit Euler) to homogenize the actual dynamics of a linear elastic composite bird with initial forcing on the wings, with (bottom) and without (top) corotational treatment. Tracking the vertical component of the linear deformation of a wing tip (indicated as a green dot on the first column) for both a fine (blue curve) and a coarsened (red curve) simulation shows a small time shift growing linearly over time, moderately increased if sparsification is introduced and only slightly affected by our RS correction in the corotational case.

6 CONCLUSIONS

Our hierarchical construction of material-adapted, matrix-based basis functions offers a new approach for coarse-graining of complex inhomogeneous materials. While we showed that a linear elastic model can be coarse-grained and post-warped through two back-substitutions per frame, refinable basis functions that are localized in space and eigenspace offer a rich framework that may be further exploited in the context of large-scale simulation. For instance, they offer an intriguing extension to the CHARMS framework [Grinspun et al. 2002], for which the decay of wavelet magnitudes should be near optimal. Its use to potentially improve geometric multigrid

methods for rough coefficients as discussed early on is also worth exploring thoroughly.

Future work on deformable models. Extending our work to non-linear elastic models is high on our priority list. An obvious approach would be to reconstruct the material-adapted hierarchy at each time step based on the linearization of the model around the current shape. Such an approach would, however, require carefully factorizing the hierarchical construction to speed it up, possibly through approximations of the current variational definition. While matching the theoretical complexity in $O(n_q \log^{2d+1} n_q)$ of the construction in the scalar case (see Fig. 13), our current implementation of the vector-valued case is not fast enough to offer much computational gain over fine simulation if the construction needs to be reevaluated every few time steps. It could be interesting as well to explore the idea of using wavelets only near material changes, an idea known as wavelet refinement in the literature, to further improve the computational gain of our approach. It may also be interesting to see if the use of non-linear potentials for which the stiffness matrix is nearly constant and for which one can simply add non-linear forces [Huang et al. 2006] could be a good application of our approach, as they would require a single hierarchy construction as preprocessing; similarly, the current trend of using projective dynamics [Bouaziz et al. 2014] could benefit from our framework. The fact that we can now homogenize an integrator step is also something that we believe deserves closer attention, to understand how the temporal and spatial homogenizations interact depending on the integrator we chose or the size of the time step—we have not looked at this at all in this work, and the topic of dynamic homogenization is certainly far from mature. Finally, our construction of elasticity-aware wavelets may also allow the injection of secondary effects in the coarse motion to add visually-enriching detail without having to simulate these high frequencies—tantamount to the way wavelets are used in fluid simulation to add visual complexity to low-resolution simulation [Kim et al. 2008].

ACKNOWLEDGMENTS

This work was supported by National Key R&D Program of China (No. 2017YFB1002703), NSFC (No. 61522209). MD acknowledges generous support from Pixar Animation Studios, and early discussion with Rasmus Tamstorf. Finally, we thank Santiago V. Lombeyda for the teaser, Tianyu Wang for generating elastic materials, and Chongyao Zhao for help with rendering.

REFERENCES

- Raymond Alcouffe, Achi Brandt, Joel Dendy, Jr., and James W. Painter. 1981. The Multi-Grid Method for the Diffusion Equation with Strongly Discontinuous Coefficients. *SIAM J. Sci. Stat. Comput.* 2, 4 (1981), 430–454.
- Ivo Babuška and John E. Osborn. 2000. Can a Finite Element Method Perform Arbitrarily Badly? *Math. Comp.* 69, 230 (2000), 443–462.
- Jernej Barbič and Doug L. James. 2005. Real-Time Subspace Integration for St. Venant-Kirchhoff Deformable Models. *ACM Trans. Graph.* 24, 3 (July 2005), 982–990.
- Jernej Barbič and Jovan Popović. 2008. Real-time Control of Physically Based Simulations Using Gentle Forces. In *ACM SIGGRAPH Asia Proceedings*. Article 163.
- Sofien Bouaziz, Sebastian Martin, Tiantian Liu, Ladislav Kavan, and Mark Pauly. 2014. Projective Dynamics: Fusing Constraint Projections for Fast Simulation. *ACM Trans. Graph.* 33, 4, Article 154 (2014).
- Achi Brandt, James Brannick, Karsten Kahl, and Irene Livshits. 2011. Bootstrap AMG. *SIAM J. Scientific Computing* 33, 2 (2011), 612–632.

- Christopher Brandt, Elmar Eisemann, and Klaus Hildebrandt. 2018. Hyper-reduced projective dynamics. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 80.
- Christopher Brandt and Klaus Hildebrandt. 2017. Compressed vibration modes of elastic bodies. *Computer Aided Geometric Design* 52–53 (2017), 297–312.
- Mary E. Brewster and Gregory Beylkin. 1995. A Multiresolution Strategy for Numerical Homogenization. *Applied and Computational Harmonic Analysis* 2, 4 (1995), 327–349.
- Max Budninskiy, Houman Owahdi, and Mathieu Desbrun. 2019. Operator-adapted wavelets for finite-element differential forms. *J. Comp. Phys.* 388 (2019), 144–177.
- Desai Chen, David IW Levin, Wojciech Matusik, and Danny M Kaufman. 2017. Dynamics-aware numerical coarsening for fabrication design. *ACM Trans. Graph.* 36, 4, Article 84 (2017).
- Desai Chen, David IW Levin, Shinjiro Sueda, and Wojciech Matusik. 2015. Data-driven finite elements for geometry and material design. *ACM Trans. Graph.* 34, 4, Article 74 (2015).
- Jiong Chen, Hujun Bao, Tianyu Wang, Mathieu Desbrun, and Jin Huang. 2018. Numerical Coarsening Using Discontinuous Shape Functions. *ACM Trans. Graph.* 37, 4 (July 2018), Art. 120.
- Yuju Chen, David I. W. Levin, Danny Kaufmann, Uri Ascher, and Dinesh K. Pai. 2019. EigenFit for Consistent Elastodynamic Simulation Across Mesh Resolution. In *Symp. Comp. Anim.* Art. 5.
- Min Gyu Choi and Hyeong-Seok Ko. 2005. Modal warping: real-time simulation of large rotational deformation and manipulation. *IEEE Trans. Vis. Comp. Graph.* 11, 1 (2005), 91–101.
- Timothy A. Davis. 2009. CHOLMOD: a sparse Cholesky factorization and modification package. <http://faculty.cse.tamu.edu/davis/suitesparse.html>.
- Tyler De Witt, Christian Lessig, and Eugene Fiume. 2012. Fluid Simulation Using Laplacian Eigenfunctions. *ACM Trans. Graph.* 31, 1 (2012), Art. 10.
- Gilles Debunne, Mathieu Desbrun, Marie-Paule Cani, and Alan H. Barr. 2001. Dynamic Real-time Deformations Using Space & Time Adaptive Sampling. In *ACM SIGGRAPH Proceedings*. 31–36.
- Christian Dick, Joachim Georgii, and Ruediger Westermann. 2011. A Hexahedral Multigrid Approach for Simulating Cuts in Deformable Objects. *IEEE Trans. Vis. Comp. Graph.* 17, 11 (2011), 1663–1675.
- Mihai Dorobantu and Bjorn Engquist. 1998. Wavelet-Based Numerical Homogenization. *SIAM J. Numer. Anal.* 35, 2 (1998), 540–559.
- Joachim Georgii and Rüdiger Westermann. 2008. Corotated Finite Elements Made Fast and Stable. In *Workshop in Virtual Reality Interactions and Physical Simulation*.
- Eitan Grinspun, Petr Krysl, and Peter Schröder. 2002. CHARMS: A Simple Framework for Adaptive Simulation. *ACM Trans. Graph.* 21, 3 (2002), 281–290.
- Robert L. Harder and Robert N. Desmarais. 1972. Interpolation using surface splines. *J. Aircraft* 9 (1972), 189–191.
- Kris K. Hauser, Chen Shen, and James F. O'Brien. 2003. Interactive Deformation Using Modal Analysis with Constraints. *Graphics Interface* (2003), 247–256.
- Michael Hautz and Wolfgang Strasser. 2004. Corotational simulation of deformable solids. *Journal of WSCG* 12, 1–3 (2004), 1–8.
- Klaus Hildebrandt, Christian Schulz, Christoph von Tycowicz, and Konrad Polthier. 2012. Interactive Spacetime Control of Deformable Objects. *ACM Trans. Graph.* 31, 4, Article 71 (2012).
- Jin Huang, Xiaohan Shi, Xinguo Liu, Kun Zhou, Baining Guo, and Hujun Bao. 2006. Geometrically based potential energy for simulating deformable objects. *The Visual Computer* 22, 9 (2006), 740–748.
- Jin Huang, Yiyong Tong, Kun Zhou, Hujun Bao, and Mathieu Desbrun. 2011. Interactive Shape Interpolation Through Controllable Dynamic Deformation. *IEEE Trans. Vis. Comp. Graph.* 17, 7 (2011), 983–992.
- Geoffrey Irving, Joseph Teran, and Ron Fedkiw. 2006. Tetrahedral and hexahedral invertible finite elements. *Graphical Models* 68, 2 (2006), 66–89.
- Inyong Jeon, Kwang-Jin Choi, Tae-Yong Kim, Bong-Ouk Choi, and Hyeong-Seok Ko. 2013. Constraining multigrid for cloth. *Comp. Graph. Forum* 32, 7 (2013), 31–39.
- Lily Kharevych, Patrick Mullen, Houman Owahdi, and Mathieu Desbrun. 2009. Numerical coarsening of inhomogeneous elastic materials. *ACM Trans. Graph.* 28, 3, Article 51 (2009).
- Theodore Kim, Nils Thürey, Doug James, and Markus Gross. 2008. Wavelet Turbulence for Fluid Simulation. In *ACM SIGGRAPH Proceedings*. Article 50.
- Marin Kobilarov, Keenan Crane, and Mathieu Desbrun. 2009. Lie Group Integrators for Animation and Control of Vehicles. *ACM Trans. Graph.* 28, 2, Article 16 (2009).
- Peter Krysl, Sanjay Lall, and Jerrold E. Marsden. 2001. Dimensional model reduction in non-linear finite element dynamics of solids and structures. *Int. J. Num. Methods Eng.* 51, 4 (2001), 479–504.
- Tassilo Kugelstadt, Dan Koschier, and Jan Bender. 2018. Fast Corotated FEM using Operator Splitting. *Comp. Graph. Forum* 37, 8 (2018), 149–160.
- Siwang Li, Jin Huang, Fernando de Goes, Xiaogang Jin, Hujun Bao, and Mathieu Desbrun. 2014. Space-time Editing of Elastic Motion Through Material Optimization and Reduction. *ACM Trans. Graph.* 33, 4, Article 108 (2014).
- Beibei Liu, Gemma Mason, Julian Hodgson, Yiyong Tong, and Mathieu Desbrun. 2015. Model-reduced Variational Fluid Simulation. *ACM Trans. Graph.* 34, 6, Article 244 (2015).
- Hsueh-Ti Derek Liu, Alec Jacobson, and Maks Ovsjanikov. 2019. Spectral Coarsening of Geometric Operators. *ACM Trans. Graph.* 38, 4, Article 105 (July 2019), 13 pages. <https://doi.org/10.1145/3306346.3322953>
- Stéphane Mallat. 2008. *A Wavelet Tour of Signal Processing: the Sparse Way*. Academic Press.
- Sebastian Martin, Bernhard Thomaszewski, Eitan Grinspun, and Markus Gross. 2011. Example-based elastic materials. *ACM Trans. Graph.* 30, 4, Article 72 (2011).
- Aleka McAdams, Yongning Zhu, Andrew Selle, Mark Empey, Rasmus Tamstorf, Joseph Teran, and Eftychios Sifakis. 2011. Efficient elasticity for character skinning with contact and collisions. *ACM Trans. Graph.* 30, 4 (2011), 37.
- Charles A. Micchelli and Theodore J. Rivlin. 1977. A survey of optimal recovery. In *Optimal Estimation in Approximation Theory*. 1–54.
- Matthias Müller, Julie Dorsey, Leonard McMillan, Robert Jagnow, and Barbara Cutler. 2002. Stable Real-time Deformations. In *Symp. Comp. Anim.* 49–54.
- Rahul Narain, Armin Samii, and James F. O'Brien. 2012. Adaptive Anisotropic Remeshing for Cloth Simulation. *ACM Trans. Graph.* 31, 6, Article 152 (2012).
- Matthieu Nesme, Paul G Kry, Lenka Jeřábková, and François Faure. 2009. Preserving topology and elasticity for embedded deformable models. *ACM Trans. Graph.* 28, 3, Article 52 (2009).
- Houman Owahdi. 2017. Multigrid with rough coefficients and multiresolution operator decomposition from hierarchical information games. *SIAM Rev.* 59, 1 (2017), 99–149.
- Houman Owahdi and Lei Zhang. 2017. Gamblets for opening the complexity-bottleneck of implicit schemes for hyperbolic and parabolic ODEs/PDEs with rough coefficients. *J. Comp. Phys.* 347 (2017), 99–128.
- Zherong Pan, Hujun Bao, and Jin Huang. 2015. Subspace dynamic simulation using rotation-strain coordinates. *ACM Trans. Graph.* 34, 6, Article 242 (2015).
- Julian Panetta, Qingnan Zhou, Luigi Malomo, Nico Pietroni, Paolo Cignoni, and Denis Zorin. 2015. Elastic Textures for Additive Fabrication. *ACM Trans. Graph.* 34, 4, Article 135 (2015).
- Alex Pentland and John Williams. 1989. Good Vibrations: Modal Dynamics for Graphics and Animation. *SIGGRAPH Comput. Graph.* 23, 3 (1989), 207–214.
- Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. 2015. Microstructures to control elasticity in 3D printing. *ACM Trans. Graph.* 34, 4, Article 136 (2015).
- Christian Schumacher, Steve Marschner, Markus Cross, and Bernhard Thomaszewski. 2018. Mechanical characterization of structured sheet materials. *ACM Trans. Graph.* 37, 4 (2018), 148.
- Eric J. Stollnitz, Tony D. Deroese, and David H. Salesin. 1996. *Wavelets for Computer Graphics: Theory and Applications*. Morgan Kaufmann Publishers Inc.
- Raghunathan Sudarshan. 2005. *Operator-adapted Finite Element Wavelets: theory and applications to a posteriori error estimation and adaptive computational modeling*. Ph.D. Dissertation. Department of Civil and Environmental Engineering, Massachusetts Institute of Technology.
- Rasmus Tamstorf, Toby Jones, and Stephen F. McCormick. 2015. Smoothed aggregation multigrid for cloth simulation. *ACM Trans Graph.* 34, 6 (2015), Art. 245.
- Rosell Torres, Jose M. Espadero, Felipe A. Calvo, and Miguel A. Otaduy. 2014. Interactive Deformation of Heterogeneous Volume Data. *Lecture Notes in Computer Science* 8789 (2014).
- Adrien Treuille, Andrew Lewis, and Zoran Popović. 2006. Model Reduction for Real-time Fluids. In *ACM SIGGRAPH Proceedings*. 826–834.
- Panayot S. Vassilevski. 2010. General constrained energy minimization interpolation mappings for AMG. *SIAM J. Sci. Comput.* 32, 1 (2010), 1–13.
- Gregory H Wannier. 1937. The structure of electronic excitation levels in insulating crystals. *Physical Review* 52, 3 (1937), 191.
- Irad Yavneh. 2006. Why Multigrid Methods Are So Efficient. *Computing in Science Engineering* 8, 6 (2006), 12–22.
- Yongning Zhu, Eftychios Sifakis, Joseph Teran, and Achi Brandt. 2010. An efficient multigrid method for the simulation of high-resolution elastic solids. *ACM Trans. Graph.* 29, 2 (2010), 16.