

Numerical Coarsening using Discontinuous Shape Functions

JIONG CHEN, State Key Lab of CAD&CG, Zhejiang University
HUJUN BAO*, State Key Lab of CAD&CG, Zhejiang University
TIANYU WANG, State Key Lab of CAD&CG, Zhejiang University
MATHIEU DESBRUN, Caltech
JIN HUANG*, State Key Lab of CAD&CG, Zhejiang University




Fig. 1. **Numerical coarsening of nonlinear materials.** We introduce a new approach for numerically capturing the behavior of inhomogeneous and nonlinear materials on coarse grids. For two composite models (left: plane with downwards forces exerted on wings; right: vase being twisted) made of two different compressible neo-Hookean materials discretized on a fine mesh (Poisson ratio: 0.45; Young's moduli: $1e3$ and $5e4$; volume rendering insets illustrate the spatial distribution of their composite material), our coarsened model with eight times less elements (red) perfectly reproduces the original behavior, while a regular FEM simulation using trilinear shape functions (green) on the same coarse mesh leads to drastic overstiffening.

In this paper, an efficient and scalable approach for simulating inhomogeneous and non-linear elastic objects is introduced. Our numerical coarsening approach consists in optimizing non-conforming and matrix-valued shape functions to allow for predictive simulation of heterogeneous materials with non-linear constitutive laws even on coarse grids, thus saving orders of magnitude in computational time compared to traditional finite element computations. The set of local shape functions over coarse elements is carefully tailored in a preprocessing step to balance geometric continuity and local material stiffness. In particular, we do not impose continuity of our material-aware shape functions between neighboring elements to significantly reduce the fictitious numerical stiffness that conforming bases induce; however, we enforce crucial geometric and physical properties such as partition of unity and exact reproduction of representative fine displacements to eschew the use of discontinuous Galerkin methods. We demonstrate that we can simulate, with no parameter tuning, inhomogeneous and non-linear materials significantly better than previous approaches that traditionally try to homogenize the constitutive model instead.

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

Additional Key Words and Phrases: Numerical coarsening, nonlinear constitutive models, deformable body simulation.

*Corresponding authors: {bao, hj}@cad.zju.edu.cn.
Authors' addresses: J. Chen, H. Bao, T. Wang, J. Huang, Zhejiang University, China; M. Desbrun, California Institute of Technology, 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.

© 2018 Association for Computing Machinery.
0730-0301/2018/8-ART120 \$15.00
<https://doi.org/10.1145/3197517.3201386>

ACM Reference Format:

Jiong Chen, Hujun Bao*, Tianyu Wang, Mathieu Desbrun, and Jin Huang*. 2018. Numerical Coarsening using Discontinuous Shape Functions. *ACM Trans. Graph.* 37, 4, Article 120 (August 2018), 12 pages. <https://doi.org/10.1145/3197517.3201386>

1 INTRODUCTION

Efficacy in simulating complex deformable models is a long standing goal in computer animation. As the demand for ever larger and more complex simulation increases, so does the need for more elaborate numerical methods whose computational costs grows only slowly with the structural complexity of the object being simulated. Traditional animation methods can only capture the proper behavior of a physical object if one uses a mesh fine enough to resolve the small-scale heterogeneities, leading to prohibitive simulation time when the fine scales are geometrically complex or have a drastically different stiffness (such as veinal structures in an organ or microstructures in a metamaterial). Simply ignoring these fine scales dramatically affects the overall dynamics of the object: a coarse simulation can fail to capture even the simplest deformation, at times rendering the object dramatically more rigid than in reality—an archetypical example being the issue of locking in finite element modeling, where incompressibility is at odds with the conforming constraints of low-order finite elements. As visual and manufacturing fidelity calls for an increasingly widespread use of non-linear geometric models and non-linear constitutive laws, the design of numerical algorithms that combine efficiency and scalability is increasingly pressing in graphics and computational physics.

Over the years, different strategies have been devised to address this enduring problem in simulation. A first family of efforts uses adaptive solvers to counteract the spurious numerical rigidity and

obvious visual artifacts that coarse elements and low-order “shape” (or “basis”) functions typically generate [Debunne et al. 2001; Grinspun et al. 2002; Narain et al. 2012]. Refinements (in the number of elements and/or their polynomial orders) based on the local amount of deformation are used to distribute degrees of freedom (DOFs) where they are most needed. However, adaptive methods are often difficult to implement without exhibiting popping artifacts, and generally not very efficient at capturing the correct behavior of heterogeneous materials at low computational cost. A second family of approaches focuses on model reduction instead, where the state space dimension is kept small by limiting the space of possible deformations, typically through modal analysis [Barbič and James 2005; Krysl et al. 2001; Li et al. 2014; Pentland and Williams 1989]. Relying on sparser approximations of the deformation space of a particular homogeneous physical model can be very efficient when combined with cubature schemes. However, model-reduced approaches typically introduce a mismatch between geometric and deformation DOFs obfuscating the treatment of contacts and collisions. They also suffer from a crippling high runtime complexity. More recently, the concept of “numerical coarsening” has shown great promise. As numerical outgrowth of homogenization theory [Bensoussan et al. 1978], coarsening refers to the idea that a complex geometric model is embedded in a coarse mesh which is assigned “coarsened” (i.e., physically averaged) material properties to best match the behavior of the original object [Kharevych et al. 2009; Nesme et al. 2009; Panetta et al. 2015; Torres et al. 2016]. Coarsening has been proposed through global (mesh-wide) or local (element-wide) optimization such that simulation at runtime efficiently captures the correct dynamics even on coarse grids. However, methods of this class are mostly limited to linear constitutive models, even if a recent extension to non-linear models has demonstrated good behavior through stiffness scaling for *homogeneous* materials [Chen et al. 2017].

In this paper, we propose a new approach to numerical coarsening which captures the fine dynamics of an inhomogeneous and non-linear elastic object using only a coarse discretization of its geometry. Instead of trying to homogenize the composite material present inside each coarse element, we achieve numerical coarsening through the computation of material-aware basis functions over the coarse computational grid. We show that allowing discontinuous and matrix-based shape functions while forcing them to properly capture a set of key continuous deformations significantly reduces the important issue of *inter-element stiffness*, i.e., the spurious effects of conforming coarse elements on the resulting dynamics. Unlike previous approaches, we are able to capture the non-linear stress-strain behavior of complex materials, with no parameter tuning.

1.1 Related Work

We first review relevant work in finite element modeling and numerical coarsening to help us motivate our simulation approach and contrast it with current techniques.

Finite-element based simulation. Real-life objects are often highly heterogeneous and best modeled by non-linear elastic models. Their accurate simulation can be achieved using finite element modeling (FEM) on tetrahedral or hexahedral elements, but at high computational cost as both geometric and material non-linearities require

dense discretization. Interpolation errors are also known to arise if badly-formed shape functions are employed; in particular, local gradients of these functions impose stringent requirements on simplicial elements for linear finite elements [Shewchuk 2002] and necessitate Jacobian-driven optimization for higher-order meshes [Johnen et al. 2013]. A number of approximation issues referred to as *locking* (parasitic shearing, incompressibility locking, membrane locking) may appear in FEM as well, introducing spurious over-stiffening when coarse elements are used [Prathap 1993]. While local refinements or degree elevation are typical ways to resolve them, a whole family of methods based on “discontinuous Galerkin elements” have been proposed [Cockburn et al. 2011]: by removing the constraint that elements (and thus, shape functions) must be conforming to their neighbors, the space of possible deformations is greatly increased—in fact, excessively so: it may create aphysical deformation, therefore artificial jump penalty terms have to be added to prevent large discontinuities in practice. Non-conforming elements that just share mid-edge points have been proposed for developable shells [English and Bridson 2008] instead, but their possible extension to 3D has not been validated yet. Note finally that harmonic shape functions over polyhedral elements [Martin et al. 2008] were also proposed to simulate elasticity with a greater deformation space than traditional simplicial elements; however, this approach is not well suited for inhomogeneous elastic deformation because shape functions are not adapted to the material composition. Balancing numerical efficiency and simulation accuracy thus remains an enduring difficulty of traditional finite elements, preventing important applications varying from realtime surgical training to rapid prototyping of metamaterials.

Constitutive model homogenization. Inspired by homogenization theory (which targets the effective averaging of solutions of equations with rapidly varying coefficients [Bensoussan et al. 1978]), a first set of numerical coarsening methods proposed to “upscale” the heterogeneous elastic properties represented by a fine mesh into *anisotropic* elastic properties of a coarse mesh that effectively captures the same constitutive model. Kharevych et al [Kharevych et al. 2009] compute a set of $d(d+1)/2$ *harmonic displacements* in dimension $d = 2, 3$ to capture how the fine mesh behaves when linear forces are applied to the boundary of the material; these representative deformations of the objects are just enough to compute an effective elasticity tensor for each coarse element, so that the potential energy of the coarse mesh exactly matches the integral of the potential of the fine mesh within each coarse element. Only linear elastic models can be coarsened though, limiting its use considerably (note however that it was recently used to solve the inverse problem of microstructure optimization as well [Schumacher et al. 2015], enabling the creation of metamaterials with prescribed bulk properties). Chen et al. [2015] propose a data-driven extension of this potential energy fitting idea to non-linear materials where now a coarse constitutive model is found through linear regression based on a set of deformation samples obtained from random forcing. However, each coarse element is homogenized individually, disregarding the physical coupling that neighboring elements induce. Moreover, results are highly dependent on the training set and the tuning of the parameters of the regression. While these model-fitting methods do manage to

homogenize the potential energy of complex materials, capturing the correct dynamics would also require homogenizing the kinetic energy by computing an effective mass matrix as well. In practice, however, a coarse lumped mass matrix assembled from the fine mesh is used as it introduces only minor inertial artifacts—but calibrating the frequency spectrum of homogenous materials through simple rescaling of Young’s modulus [Chen et al. 2017] has shown encouraging results towards full dynamics-aware coarsening.

Shape function optimization. Another approach for coarsening consists in computing meaningful, “material-aware” shape functions in an offline preprocessing step: using tailored basis functions on the coarse elements that closely reproduce the possible deformations of a complex material instead of the traditional polynomial functions significantly improves the accuracy of simulation. This idea was articulated as a local constrained quadratic optimization per coarse node in [Owhadi et al. 2014], while local eigenanalysis was introduced in [Efendiev et al. 2011] to tailor better material-adapted shape functions. However, the locality of the solves involved in these two methods generate shape functions that do not even satisfy partition of unity, making them not amenable to general elasticity. Nesme et al. [2009] proposed the computation of a shape function and stiffness matrix per coarse element such that the deformation computed using this single coarse element best fits the deformations of the fine mesh it represents. They used boundary vanishing conditions and assumed the interpolation for 3D displacements to be independent in each Cartesian coordinate. Torres et al. [Torres et al. 2014, 2016] further extended this approach for arbitrary boundary conditions through condensation, and proposed a discontinuous matrix-valued mapping from coarse to fine grids derived from the stiffness matrix. Note that at runtime, these methods use a corotational-based coarse simulation to improve their results by adding geometric non-linearity: it makes the effective linear strain tensor more accurate for large deformation by removing as much of the current local rotation as possible. However, none can handle fully non-linear materials.

Multigrid solvers. While fundamentally different from numerical coarsening which aims at finding a coarse solution without resorting to computations on the original fine mesh, geometric multigrid methods (GMG) use prolongation (and restriction) operators to project a coarse solution onto a fine grid (and vice versa) in order to accelerate computations. The accuracy of this mapping from coarse to fine and fine to coarse is key in the efficiency of a solver; yet currently, most multigrid methods (e.g. [McAdams et al. 2011]) simply construct prolongation operator through geometric interpolation without much consideration of the equation they are solving: any error in this mapping will get corrected based on subsequent computations on the fine grid. Our coarsening work may thus be useful for GMG as it provides an accurate finite-dimensional map between various resolutions through a careful design of shape functions.

1.2 Contributions

Based on our review of previous works, a few comments are in order. First, while it was recognized early on that an assembly of different isotropic materials is in general anisotropic [Kharevych et al. 2009], all methods so far have relied on scalar shape functions, instead of embracing the unavoidable anisotropy by using general tensorial

shape functions to interpolate through linear transformations. Second, current coarsening methods upscale each coarse element, but do not address the negative (generally over-stiffening) effects that conforming (geometrically continuous) discretizations can induce.

In this paper, we show that coarsening via the construction of *discontinuous and matrix-valued* shape functions satisfying important geometric and physical material-aware properties offers a significantly more flexible and general approach to coarsening. From a pair of regular grids representing respectively a fine and a coarse geometric description of the elastic body, we devise a numerical procedure to turn the heterogeneous and non-linear elastic properties of the fine mesh into a set of material-aware shape functions over the coarse mesh to better capture the solution space of the fine mesh. Our matrix-valued shape functions are expressed in individual corotational frames, reminiscent of the frame-based approach of [Faure et al. 2011; Gilles et al. 2011], to better capture the non-linearity of a material on the coarse grid. Optimizing these shape functions involves no heuristics or parameter tuning, but a simple constrained quadratic programming problem which adapt these functions to the material while keeping key properties such as partition of unity and infinitesimal rotation invariance. The resulting shape functions do not suffer from the fictitious numerical stiffness of conforming bases, and exactly reproduce a set of representative deformations of the object. The coarsened dynamical system can then be deformed with a computational complexity nearly proportional to the size of the coarse mesh for any non-linear constitutive model, with a visual accuracy decaying gracefully with the size of the coarse mesh.

2 DISCONTINUOUS, TENSORIAL SHAPE FUNCTIONS

We begin our exposition by defining the goals of numerical coarsening, and deriving the general form that our shape functions will take to be sufficiently expressive to capture complex deformation.

2.1 Elastic behavior

Given a deformable object occupying a d -dimensional spatial domain $\Omega \subset \mathbb{R}^d$ at rest, its deformation can be encoded via a displacement field $u : \Omega \rightarrow \mathbb{R}^d$. For a given external force field (and possibly some position constraints) the object will deform, giving rise to a deformation field over the reference (material) domain Ω . Mathematically, the resulting displacement u^* can be described variationally: it is a constrained minimization (sometimes referred to as a *elastostatics* problem), which defines u^* as the deformation leading to the minimum elastic potential energy subject to external forces and constraints; i.e.,

$$u^* = \arg \min_u \int_{\Omega} \Psi(u) dX + \int_{\Omega} \langle f, u \rangle dX \quad \text{s. t.} \quad C(u) = 0 \quad (1)$$

where Ψ is the elastic potential density of the object defined over the undeformed shape Ω , f represents the external force field, and C encodes the deformation constraints (such as prescribed nodal displacements for instance).

2.2 Fine vs. Coarsened Finite Element Modeling

Finding a deformation given external forces requires computations. Let us describe how this would be done on a fine mesh, then how this could be approximated on a coarsened mesh instead.

Fine shape functions. To numerically solve the problem given by Eq. (1), Finite Element Modeling (FEM) proceeds by discretizing the domain Ω into a fine mesh Ω^h with a set of nodes X_i^h and a number of elements Ω_e^h formed by these nodes. For each element Ω_e^h , a scalar shape function $N_{e,i}^h: \Omega \rightarrow \mathbb{R}$ per corner node X_i is defined as well so that they sum to one (partition of unity) and $N_{e,i}^h(X_j^h) = \delta_{ij}$ (Lagrange property). Linear functions for tetrahedron meshes and trilinear functions for regular grids are often selected for their simplicity. With this setup, a continuous deformation u can be encoded as one displacement $u_i^h = u(X_i^h)$ per node X_i^h , since the deformation field can be interpolated to the rest of the domain via:

$$\forall X \in \Omega_e^h, \quad u(X) = \sum_{X_i^h \in \Omega_e^h} N_{e,i}^h(X) u_i^h. \quad (2)$$

Note that the discretization presented here is far from general: many variants exist. However, this particular one is commonly used because it is very local (each node only has an influence on its immediate neighboring elements) and interpolating (due to the Lagrange property). Moreover, the variational problem of Eq. (1) can now be discretized and solved efficiently for a function u^* in a finite dimensional subspace parameterized by $\{u_i^h\}_i$.

Coarsened shape functions. For the above FEM procedure to provide an accurate deformation, each mesh element should only describe a small and homogeneous part of the deformable object. As a consequence, finding the deformation of an inhomogeneous object with high contrast in material stiffness and strong non-linear stress-strain behavior requires an inordinate count of elements, precipitously increasing the computational cost involved in dealing with such a fine mesh. Therefore, most numerical coarsening methods begin by reducing the spatial resolution, through a coarse mesh Ω^H for which each element Ω_e^H is an assembly of fine elements from Ω^h , and where the nodes of Ω^H are a subset of the fine nodes of Ω^h . We adopt the same strategy in our work: we assume that both fine and coarse meshes are *regular grids* to make our explanations (and the implementation) easier, but other spatial arrangements can be used as well. For this coarse mesh to be able to resolve nearly the same deformation at a fraction of the computational cost, we need to optimize the shape functions $N_{e,i}^H$: using the same trilinear basis functions as on the fine grid would simply fail to capture the complexity of the physical behavior of the deformable object (see Figs. 1 and 18 for instance). Instead, one should use shape functions such that their span covers the span of the fine basis functions, and is expressive enough to reproduce the typical deformation of the object at hand: only then can we hope that solving the variational problem of Eq. (1) on the coarse grid is nearly equivalent to the more computationally intensive solve on the fine grid.

For simplicity, we will omit the subscript e from now on: the notation N_i^H will be used instead of $N_{e,i}^H$ to refer to a basis function associated to the node X_i^H of a certain coarse element Ω_e^H .

2.3 Piecewise-trilinear shape functions

In order to offer enough flexibility compared to the fine grid, each coarse shape function N_i^H over a coarse element Ω_e^H for a coarse node $X_i^H \in \Omega_e^H$ can be defined using the fine grid: the shape function




Fig. 2. **Piecewise bilinear shape functions.** A 2D coarse element (left) and its associated fine elements and fine vertices used to discretize its shape functions (center). Each fine node is associated with a local basis function (right, depicted via height field), which consists of piecewise bilinear functions defined on its directly adjacent fine elements.

is discretized by values on the fine nodes X_j^h contained within this coarse element, and interpolated via the fine shape functions N_j^h (see the 2D illustration in Fig. 2). In other words, while the fine mesh uses trilinear shape functions, a coarse shape function N_i^H over the element Ω_e^H is now a piecewise-trilinear function defined as:

$$\forall X \in \Omega_e^H, \quad N_i^H(X) = \sum_j n_{ij} N_j^h(X), \quad (3)$$

where the coefficients n_{ij} are now degrees of freedom that can be optimized at will, representing the coarse nodal basis N_i^H sampled on the fine nodes $X_j^h \in \Omega_e^h$.

2.4 Corotational matrix-valued shape functions

While piecewise-trilinear coarse shape functions offer much added degrees of freedom, they suffer a key limitation: when used to interpolate a deformation from coarse values u_i^H on the coarse mesh, the deformation field will interpolate each coordinates of u_i^H *independently*, since now Eq. (2) reads on the coarse mesh as:

$$\forall X \in \Omega_e^H, \quad u(X) = \sum_{X_i^H \in \Omega_e^H} N_i^H(X) u_i^H. \quad (4)$$

This means that any displacement in the x -direction of a node value u_i^H will *only* affect the x -coordinates of the reconstructed field—a very limiting interpolation of complex deformation.

Matrix-valued shape functions. Yet, typical elastic deformation exhibits strong coupling between coordinates, for instance in material with high Poisson ratio. We thus propose to turn the usual *scalar-valued* shape functions into **matrix-valued** shape functions:

$$N_i^H: \Omega_e \rightarrow \mathbb{R}^{d \times d},$$

to best exploit the inevitably-anisotropic behavior of complex objects: our interpolation now involves general linear transformations that couple dimensions and handle anisotropy naturally. This is easily done with our previously-introduced piecewise-trilinear setup by now considering the degrees of freedom n_{ij} to be $d \times d$ matrices instead of scalar values.

Corotational formulation. The use of tensorial shape functions comes with an immediate challenge, though. A tensor requires a frame in which to express it as a matrix. The natural choice of a fixed Cartesian frame is, alas, not appropriate since the interpolation would no longer be rotation invariant: an added global rotation would change the relative displacements within elements. This issue is easily resolved by adding a notion of local frame, similar to what is widely used in corotational methods. First, we estimate a local frame

R_e for each coarse element Ω_e^H through a simple procedure: we interpolate the current node displacements via simple trilinear shape functions \bar{N}_i over the coarse element and evaluate the deformation gradient $\nabla x = \nabla u + \mathbf{I}$ at the center of the element; the frame R_e is then defined as the rotation matrix of the polar decomposition of ∇x . The shape functions are then used in this local frame to interpolate the deformation field through:

$$u(X) = R_e \left[X + \sum_i N_i^H(X) (R_e^T x_i^H - X_i^H) \right] - X. \quad (5)$$

When $N_i^H(X)$ is linear precise (i.e., $\sum_i N_i^H(X) X_i^H = X \forall X$, which holds for trilinear functions \bar{N}_i), the above equation simplifies to the classical corotational formulation used in linear elasticity animation [Müller and Gross 2004]. Our expression is thus a corotational treatment of our generalized matrix-valued shape functions, allowing us to enforce an artifact-free interpolation of deformation fields within each element.

2.5 Discussion

Note that our construction of coarse shape functions uses a number of DOFs stored on the fine grid to capture the complexity of the fine structure with only a small amount of coarse nodes, as already proposed in [Owhadi et al. 2014] for instance. However, our use of $d \times d$ matrices for n_{ij} instead of scalars uses 4 times more memory in 2D, and 9 times more in 3D: we are thus, in a way, trading memory for expressiveness. The resulting coarsened shape functions will, however, allow for simulations with similar or better visual results than their less memory-consuming counterparts as we will demonstrate in Sec. 4. Finally, this choice of extended shape functions creates a possibly discontinuous deformation field when coarse deformation vectors are interpolated. We discuss next how to restrict this high dimensional set of shape functions to guarantee that they ensure a good balance between geometric continuity of the deformation and local respect of material behavior.

3 DESIGN OF MATERIAL-AWARE SHAPE FUNCTIONS

While we introduced a new expression for coarsened shape functions, not every function in this large allowable space is a valid candidate for accurate simulation. In this section, we identify key properties that shape functions should satisfy for a given inhomogeneous, non-linearly elastic material. We then propose an optimization procedure to construct the actual shape functions. Special attention is paid to balance geometric continuity and local stiffness.

3.1 Geometric conditions

Basic geometric properties of shape functions need to be preserved in our setting to avoid aphysical behavior.

Geometric invariance. If an object is simply translated or rotated, its strain should not be changed: it would introduce spurious artifacts otherwise. In order to render our discretization translation invariant, we enforce the following condition on each coarse element Ω_e^H :

$$\sum_i N_i^H(X) = \mathbf{I}, \quad (6)$$

where \mathbf{I} is the $d \times d$ identity matrix. This is simply the traditional *partition of unity*, extended to our tensor case. Similarly, if the object

rotates with a constant angular velocity ω , we need to enforce:

$$\omega \times X = \sum_i N_i^H(X) \omega \times X_i^H, \quad \forall \omega$$

to avoid spurious strains. Using the matrix notation $[\cdot]_\times$ defined such that the cross product $a \times b$ is expressed as $[b]_\times a$, this reads:

$$[X]_\times = \sum_i N_i^H(X) [X_i^H]_\times. \quad (7)$$

Node interpolation. We also enforce the traditional Lagrange property, i.e., if δ_{ij} is the Kronecker delta function, we must have:

$$N_i^H(X_i^H) = \delta_{ij} \mathbf{I}. \quad (8)$$

This property guarantees that coarse nodes are properly interpolated, which is crucial for point constraints or collision handling.

3.2 Physical conditions

Assuming they satisfy the basic geometric properties we just described, our coarsened shape functions still have the difficult role of having to reproduce as much as possible the physical behavior captured by the fine grid. In order to render our shape functions material-aware, we compute a set of representative deformations and constrain our shape functions to reproduce those *exactly*. We follow the approach of [Kharevych et al. 2009] that proposed to solve $d(d+1)/2$ global “harmonic” displacements $\{h_{ab}\}_{ab}$ ($a, b \in \{1, \dots, d\}$) on the fine grid to characterize how the fine mesh behaves under a set of chosen boundary conditions, see Fig. 3. These deformations are found by solving the following linearized elastostatic problems with Neumann boundary conditions:

$$\begin{aligned} \nabla \cdot \sigma(h_{ab}) &= 0, & \text{inside } \Omega, \\ \sigma(h_{ab}) \cdot n &= \frac{1}{2}(e_a e_b^T + e_b e_a^T) \cdot n, & \text{on } \partial\Omega, \end{aligned} \quad (9)$$

where σ is the stress tensor and e_a refers to the unit vector in the a -th coordinate direction. Note that these displacements are found through a simple linear system even if the material constitutive model is non-linear: we consider the Hessian of the potential energy to be constant, and equal to $\text{Hess}(\Psi)$ evaluated at $u=0$ (i.e., we use the stiffness matrix of the object at rest). Indeed, a full non-linear solve would necessitate a choice of boundary traction magnitude, whereas the linearized version of [Kharevych et al. 2009] needs no such parameter: in linear elasticity, the deformation is simply proportional to the traction magnitude. This amounts to a global infinitesimal “probing” of the object by a set of linear traction fields on the boundary. As proposed in [Kharevych et al. 2009] (and more recently in [Schumacher et al. 2015]), we fix the last six degrees of freedom (translation and rotation) of Eq. (9) by fixing the zeroth and first moments, resulting in a unique solution. As we will discuss later on, these deformation fields are enough to make our shape functions material-aware, thus bypassing the need for a full non-linear treatment.

Once the characteristic deformation fields h_{ab} are found, we require our shape functions to precisely reconstruct them, that is,

$$h_{ab}(X) = \sum_i N_i^H(X) h_{ab}(X_i^H). \quad (10)$$

Note that this condition imposes 6 constraints per coarse element in 3D (3 in 2D). Our shape functions will then accurately represent any deformation that is a linear combination of harmonic displacements.

Moreover, since exact reproduction of representative displacements is guaranteed by construction, it induces a weak enforcement of *geometric continuity* of the coarsened shape functions: possible discontinuities introduced by our use of local frames are suppressed on a series of characteristic deformations, without the non-physical penalty terms often used in Discontinuous Galerkin methods; a deformation will exhibit discontinuity only if it is not in the span of our harmonic displacements.

3.3 Numerical conditioning

An important measure of numerical conditioning of a shape function is through the norm of its gradient: an element is badly shaped if its basis functions are not able to capture local gradients properly, which usually negatively affects the condition number of the stiffness matrices in linear finite elements [Shewchuk 2002]. The same conditioning requirement is valid in higher order shape functions, and significant efforts in meshing are spent to make sure the local Jacobian of a parameterization is non-degenerate [Johnen et al. 2013]. In our context, the elements are not the issue: they are regular hexahedra; but their associated shape functions have not been restricted to be well conditioned in any way so far.

If our shape functions were scalar functions, a simple functional to minimize in order to offer good conditioning would be the usual Dirichlet energy, ubiquitous in graphics: its minimization favors equidistribution of gradients throughout the domain. We propose the same functional, but now for our tensor-valued functions: we use the integral over a coarse element Ω_e of the M -weighted Frobenius norm of the gradient of our coarsened shape functions, i.e.,

$$\sum_i \int_{\Omega_e} \text{tr} \left((\nabla N_i^H)^T : M : \nabla N_i^H \right) dX \quad (11)$$

where M is a given symmetric $d \times d \times d \times d$ rank-4 tensor (and ∇N_i^H is $d \times d \times d$ tensor because we use matrix-valued functions).

Choice of M . Note that if M is chosen to be the identity tensor $M = \mathbb{I}$, this functional is simply the L_2 norm of the gradient field. Minimizing this functional will equalize the spatial distribution of the shape functions, thus rendering them appropriate for numerical simulation. Another choice of M is interesting as well in our context: if one expands the nonlinear potential energy density Ψ around the




Fig. 3. **Harmonic deformation.** Using six different linear solves, the “harmonic displacements” [Kharevych et al. 2009] of a given composite material (here, an object with alternating layers of a stiff and a soft material) are computed from a set of linear traction fields on the boundary.

rest state to the second order, one finds

$$\Psi(\mathbb{I} + \nabla u) \approx \Psi(\mathbb{I}) + \frac{\partial \Psi}{\partial F}(\mathbb{I}) : \nabla u + \frac{1}{2} (\nabla u)^T : \frac{\partial^2 \Psi}{\partial F^2}(\mathbb{I}) : \nabla u,$$

where we used $F = \mathbb{I} + \nabla u$ for the deformation gradient. Evaluated at the rest state ($u = 0$), this Taylor expansion reduces to a quadratic form using the stiffness tensor $\partial^2 \Psi / \partial F^2$ at rest, already leveraged to find harmonic displacements in Section 3.2. Thus, one can also use $M = \partial^2 \Psi / \partial F^2|_{u=0}$: this is another smoothness measure of the gradient fields of our shape functions, this time adapted to the material model. We call these two conditioning functionals *harmonic* and *Ψ -harmonic* respectively.




Fig. 4. **Neo-Hookean composite bar.** For the deformation of a composite compressible neo-Hookean material (Poisson’s ratio is 0.45, Young’s modulus is 10^3 vs. $5 \cdot 10^4$) under gravity (left), using Ψ -harmonic conditioning (right) makes the shape functions less conforming, resulting in a slightly softer behavior compared to the basic harmonic conditioning (center). Bottom insets visualize the reconstructed fine mesh to exhibit more clearly the fine scale differences.

3.4 Constrained optimization

Solving for the optimal shape functions is now simple: we find, for each coarse element, the corner shape functions that satisfy all the geometric and physical conditions while being best conditioned. This is achieved through a constrained quadratic optimization on the degrees of freedom $\{n_{ij}\}_{ij}$. Thus, for each coarse element Ω_e^H , we compute the solution to the following constrained minimization:

$$\begin{aligned} \arg \min_{\{n_{ij}\}} & \sum_i \int_{\Omega_e} \text{tr} \left((\nabla N_i^H)^T : M : \nabla N_i^H \right) dX, \\ \text{s. t.} & \sum_i n_{ij} = \mathbb{I}, \quad \forall X_j^h, \\ & \sum_i n_{ij} [X_i^H]_{\times} = [X_j^h]_{\times}, \quad \forall X_j^h \\ & \sum_i n_{ij} h_{ab}(X_i^H) = h_{ab}(X_j^h), \quad \forall X_j^h, \quad \forall a, b \\ & n_{ij} = \delta_{ij} \mathbb{I}, \quad \forall X_i^H, X_j^H. \end{aligned} \quad (12)$$

This numerical optimization reduces to a quadratic optimization on the matrices n_{ij} with the linear constraints presented in Eqs. (6-8) and (10). By assembling all the DOFs of a given element into a vector n_e , the constraints can be expressed as a linear system

$$C n_e = y. \quad (13)$$

To accelerate the solve while ensuring that these hard constraints are met, we compute the kernel U of C and find an arbitrary solution

n_e^0 of Eq. (13). We then rewrite the unknown vector as

$$n_e = Uq + n_e^0, \quad (14)$$

where now the subspace coordinate q is the new, smaller vector to optimize for. In practice, we use Suitesparse’s QR solver [Davis 2017] to find the arbitrary solution and the kernel, and the reduced coordinates of the shape functions are obtained by simply solving a linear system to minimize the additional regularization. By constructing our shape functions with this optimization, the resulting coarsening procedure applies to a much wider range of hyperelastic materials, such as neo-Hookean models. This is in sharp contrast with previous work like [Kharevych et al. 2009; Nesme et al. 2009] that are limited to linear constitutive models. Moreover, because our constraints include exact preservation of key “harmonic” deformations, the numerical effect of inter-element coarsening is better accounted for as our results will demonstrate in Section 4. Finally, note that given that we use regular grids, one-level coarsening (i.e., each dimension reduced by a factor of two) results in 1944 DOFs for basis discretizations per coarse element (8 vertices, each having 27 fine nodes to store the basis, and each node storing the 9 coefficients of a matrix) while it goes up to 9000 for two-level coarsening. The kernel dimensions of the constraints (i.e., the effective DOFs) in these two cases are 684 and 4212 respectively.




Fig. 5. **Harmonic vs. Ψ -harmonic conditioning.** On a neo-Hookean composite material (left), harmonic conditioning (center) outperforms material-aware conditioning (right), although proper averaging of the fine node positions to post-process discontinuities across coarse elements leads to very similar visual results.

3.5 Balancing function continuity and local stiffness

Now that the precise optimization used to construct shape functions have been formulated, the choice of harmonic vs. Ψ -harmonic conditioning can be discussed in more detail. Ideally, basis functions should be continuous between elements *and* they should properly capture the local stiffness of their elements. But both conditions can only be obtained for very fine grids that capture all the details of a composite material.

Selecting $M = \mathbb{I}$ does promote inter-element continuity as illustrated in Fig. 4 as the local shape functions are made as smooth as possible; however, this smoothness may introduce more potential energy in coarse elements with an inhomogeneous material: forcing continuity across elements puts a limit on how well the potential energy is approximated. The Ψ -harmonic conditioning, instead, allows each single coarse element to relax to its minimal energy state while *only* keeping the common nodes between neighboring elements at the same location (due to the interpolation constraints). While this will better match the local stiffness of the fine elements, the inter-element discontinuities are allowed to be larger (see Fig. 4),

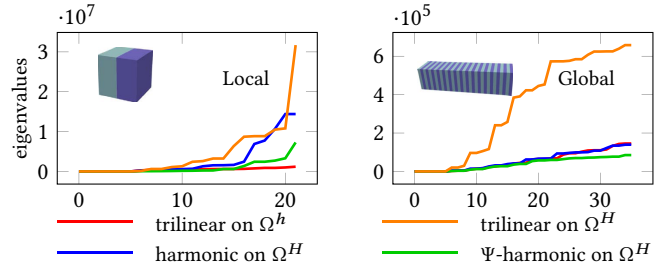


Fig. 6. **Coarsened eigenvalues.** For a two-material linear elastic bar with alternating stiffness, we compare the spectrum (eigenvalues of $Ku = \lambda Mu$, M being the lumped mass matrix) of the local stiffness matrix (for one composite coarse element, left) and of the global stiffness matrix (right). Our coarsening reproduces the spectrum significantly more accurately than trilinear basis functions.

especially if neighboring coarse elements have widely different material composition. Thus, this choice may soften the global behavior.

We will see in Section 4 that this analysis is confirmed by our experiments. Fig. 6 demonstrates this behavior by exhibiting the spectrum of the stiffness matrix of a composite linear elastic material, for a single element as well as for the whole object. In particular, it clearly highlights that traditional linear shape functions on a coarse mesh make the results clearly too stiff because it completely eliminates the inter-element discontinuity, significantly increasing the inter-element stiffness. Note again that our approach sharply contrasts with the use of jump penalty terms in Discontinuous Galerkin methods [Cockburn et al. 2011] which help to limit discontinuities through non-physical regularization.

4 RESULTS AND EVALUATION

In this section, we first discuss how our method actually simulates a coarsened object by providing a few relevant implementation details. We then demonstrate the balance between geometric continuity and material awareness that our numerical coarsening offers by comparing its advantages over other existing methods.

4.1 Simulation of Coarsened Model

Once the coarsened basis functions N_i^H have been computed for a given fine object, simulating the deformation of this object subject to an external force field and/or positional constraints proceeds mostly like with a regular Finite Element Analysis solver—with the exception that our tailored shape functions are used in lieu of the usual polynomial functions. We however point out a few implementation details that we found useful in our own implementation.

Local frame estimation. The polar decomposition of the deformation gradient at the center of the parametric domain Ω_e required to express the displacement field in Eq. (5) can be computed only once for every iteration of the solver: considering this frame constant per iteration does not give rise to convergence issues or instabilities.

Pointwise deformation gradient. Evaluating the deformation gradient can be achieved through quadrature, as typically done when dealing with higher-order polynomial shape functions. To integrate the elastic potential using a standard quadrature scheme, we map




Fig. 7. **Comparing with [Kharevych et al. 2009]**. We compare our results (red) to theirs (green) for 4 examples involving linear elastic composite materials (left). Histograms of mean and max displacement errors w.r.t. fine simulation averaged over 388 different tests are shown for both methods (right).

the element domain Ω_e to the isoparametric domain $[-1, 1]^d$ (with coordinates referred to as ξ), and apply the chain rule to evaluate the composition [Irving et al. 2006]: for each quadrature point on coarse element, the actual pointwise gradient $\nabla_X x$ of the deformation is

$$\begin{aligned} \nabla_X x &= \nabla_X u + \mathbb{I} = (R_e - \mathbb{I}) + \sum_i R_e \otimes (R_e^T x_i - X_i) : \frac{\partial N_i^H}{\partial X} + \mathbb{I} \\ &= R_e + \left(\sum_i R_e \otimes (R_e^T x_i - X_i) : \frac{\partial N_i^H}{\partial \xi} \right) \left(\sum_j X_j \frac{\partial \bar{N}_j^H}{\partial \xi} \right)^{-1} \end{aligned} \quad (15)$$

where u is the displacement field, X is the material point and \bar{N}_i^H are the conventional trilinear shape functions used in Sec. 2.4 to estimate the local frames R_e . Note that the tensor product symbol \otimes between the matrix R_e and the vector $R_e^T x_i - X_i$ generates a rank-3 tensor, which is then contracted against $\partial N_i^H / \partial X$.

Quadrature scheme. In our implementation, if a coarse element contains n^d fine elements, we use a n -point Gaussian quadrature for each dimension. The inset shows a 2D illustration for a coarse element with 4×4 fine elements. For each quadrature point, we can easily evaluate the deformation gradient using Eq. (15), from which the pointwise elastic potential and its gradient can be computed. Forces are then computed by a weighted sum of these values at quadrature points with associated weights (approximating their integral counterparts) as done in, e.g., [Irving et al. 2006].

Fine detail reconstruction. Because we enforced that our shape functions interpolate the coarse nodes, visualizing the deformed coarse grid is trivially achieved (we just displace a node X_i^H by u_i) and does not require any blending. However, our optimized shape functions also allow us to efficiently reconstruct the fine deformation that a coarse deformation captures simply by using Eq. (5): detailed anisotropic and non-uniform deformations are well captured, as shown in Fig. 8. If watertight fine meshes are required for visualization purposes, averaging of the fine nodes on each side of coarse elements can be done to remove the generally small discontinuities between elements during simulation.




Fig. 8. **Coarsened vs. fine embeddings.** The fine-scale accordion effect created by a two-material composite (St Venant-Kirchhoff) object (dark blue: stiff; light blue: soft) is well captured by the embedded 64-times finer grid (see closeup), making the coarsened simulation material-aware.

Enriching representative displacements. In order to further reduce inter-element discontinuity, a natural idea is to generate *more* representative displacements and force their exact reproduction as we did for the harmonic displacements. We tried to use the 200 leading eigenvectors of the tangent stiffness matrix $\partial^2 \Psi / \partial X^2$ at rest as additional constraints in Eq. (12). The results were nearly identical, implying that this large number of additional constraints were, in fact, nearly colinear to the six harmonic ones per coarse element. Using global harmonic displacement or eigenvectors at *deformed* states in a principled manner (and not just through random forcing which could stiffen the results significantly) would be an interesting extension that we leave as future work.

4.2 Advantages over representative methods

We tested our approach on a number of examples to demonstrate its advantages compared to basic finite element analysis for coarse grids and existing coarsening methods. Three different elastic models are employed in this paper: linear elasticity, St Venant-Kirchhoff models, and neo-Hookean models. For each composite model, Young's modulus of the stiffer material is 50 times the one of the softer material. The Poisson's ratio of all the materials is set to 0.45 (except for the stress test in Fig. 9 where we use 0.4999 to illustrate the absence of locking or overstiffening in this challenging case).

Coarsening of homogeneous material. As a first test, we try our method on a *homogeneous* material. As shown in Fig. 10, our results obviously outperform the typical use of linear shape functions on such a coarse grid: while usual FEM elements lead to significant over-stiffening on coarse resolutions due to low-order polynomial functions and conforming elements, our discontinuous tensorial shape functions allow for a very accurate capture of deformation.




Fig. 9. **High Poisson’s ratio.** Coarsening (center) of a homogeneous neo-Hookean material with Poisson’s ratio 0.4999 shows nearly perfect agreement with the fine mesh simulation (left), while a traditional FEM simulation (right) fails to handle this challenging case.




Fig. 10. **Neo-Hookean arch.** An arch-shaped bar made out of a homogeneous neo-Hookean material is bent via a fine grid (left); our coarsened simulation matches the shape (center); regular trilinear shape functions fail to capture large deformation in the legs and the curvature of the bar (right).

Relation to [Kharevych et al. 2009]. The most significant advantage of our method over the work of Kharevych et al. [2009] is the ability to handle non-linear constitutive models (e.g., the neo-Hookean material in Fig. 4). If we apply our method to corotational linear elasticity, our results are visually comparable to their upscaling of the elasticity tensor, which is to be expected as we use their harmonic deformations to optimize our shape functions. We even perform better in terms of coarse node displacement errors (both in mean and max) as illustrated by the histogram in Fig. 7, obtained by running 388 different deformation test examples.

Relation to [Nesme et al. 2009]. Similarly, the original approach of Nesme et al. [2009] cannot handle non-linear materials either. While they also use a co-rotational treatment, our use of tensorial shape functions offers more degrees of freedom to capture the behavior of an arbitrary composite material. Moreover, one can argue that their use of a clamped far boundary condition per element helps improve inter-element continuity by construction. The obvious adverse effect is, however, an increase in the global stiffness. We can try to reproduce this condition in our method by adding it as a constraint—at the cost of removing the constraint of exact preservation of the key “harmonic” deformations (see Section 3.2) since the two constraints would be in conflict. For numerical conditioning, the Ψ -harmonic regularization is better than the harmonic one in this case ($M = \mathbb{I}$ is nearly identical to using trilinear functions). Still, as shown in Fig. 11, this clamped boundary treatment induces stiffness by favoring conforming elements, indicating that their per-coordinate coarsening approach with clamped boundary fails to capture complex material behavior. Our results are much closer to the real material’s behavior, as our matrix-valued shape functions reproducing harmonic displacements reduce the local stiffness engendered by neighboring elements as much as possible.

Relation to DDFEM. Our experiments with [Chen et al. 2015] indicates a strong dependence of their technique on the training set: following their dataset generation strategy, we generated two different training sets using different magnitudes of forces, and find the results to be significantly different, see Fig. 12. Their reliance on a regularization weight also requires time-consuming manual tuning. If a small regularization is chosen, their per-element coarsening with no conforming constraints leads to overly soft results; conversely, a high regularization weight will overly stiffen all coarse elements, see Fig. 13. This is a serious practical limitation: if different parts of an object have different material compositions (e.g., a bar with its left part being homogeneous and its right part being inhomogeneous), no uniform weighting parameter will be adequate: a proper weighting really depends on the material composition in the neighboring elements. With our global harmonic reconstruction constraints, our method successfully gets rid of the need for a manually fine-tuned parameter or a choice of force magnitude.

4.3 Conditioning choice

We also tested both conditioning functionals discussed in Section 3.3. The material-based conditioning induces softer results than the fine mesh because inter-element discontinuity is not as strongly penalized. Instead, the identity metric usually leads to better balance between intra-element material awareness and inter-element conformation. We provide more comparisons between these two metrics in Fig. 5. While the two functionals can be used, we recommend the use of the simple unweighted Frobenius norm in practice.




Fig. 11. **Deformation of neo-Hookean models.** On four different tests, a fine mesh simulation (left) is well reproduced by our coarsened grid where harmonic conditioning has been used (center). For comparison, we also show a coarsened model where clamped far boundary conditions [Nesme et al. 2009] are imposed in lieu of harmonic displacements (right).




Fig. 12. **Influence of training set on DDFEM.** The DDFEM approach [Chen et al. 2015] depends heavily on the training set, even for homogeneous materials: a set of deformations based on small magnitude random forcing (left) generates very different results than with 100-times stronger forcing (right). Since force fields with larger variations will lead to extreme local deformation, the incompatible energy profile makes the results much softer. A neo-Hookean model is used here.

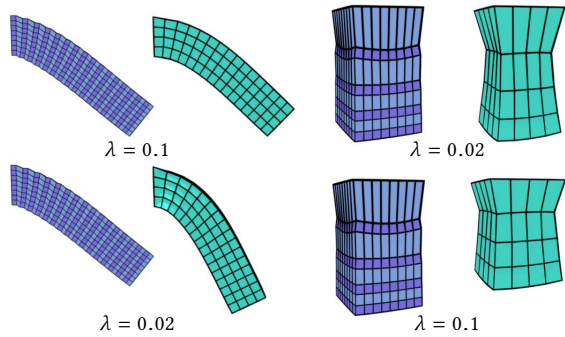


Fig. 13. **Limitations of DDFEM.** Although these neo-Hookean objects are composed of the exact same two materials sandwiched in the same pattern, one has to apply different regularization weights to get visually consistent results on coarse and fine meshes (upper row). Different weights leads to significant error (bottom row).

4.4 Large deformation

Even though our method gets much lower mean displacement errors than with traditional trilinear basis functions, the approximation error of a coarse discretization increases with the magnitude of the deformation. Consequently, our optimized basis functions may fail to capture very large and highly nonlinear deformation, as shown in Fig. 14: after all, the basis functions are only capable of exactly reproducing the representative displacements at the rest state. Since the basis functions are locally supported, we can still expect low approximation error for a range of deformation even if the basis functions are optimized around the rest state. However, very large deformation of highly nonlinear elastic composite models are bound to be challenging for our coarsening approach. How to further transform basis functions to properly adapt to very large deformation is an important and challenging problem to be tackled.

4.5 Coarsening of dynamics

Since our homogenized shape functions can capture a broad range of deformation, they are ideally suited for elastostatics, i.e., finding deformation for given force fields. However, they can also be used in the more general context of dynamics using a lumped mass matrix for coarse elements that is simply constructed by summing the contributions from the fine mass matrix (see Figs. 15 and 16). This is, however, an approximation as recognized in previous work [Chen

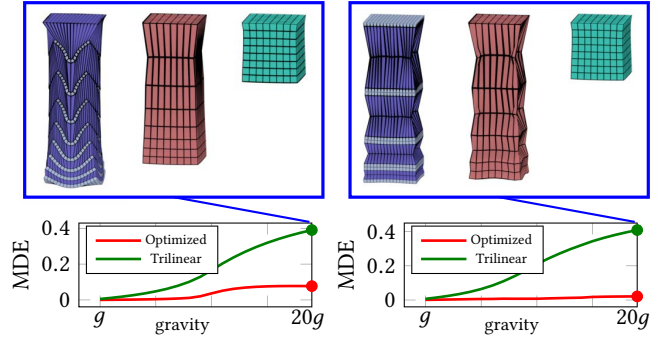


Fig. 14. **Displacement error for large deformation.** Under various gravitational accelerations (from g to $20g$), deformation grows and so does the mean displacement error (MDE) for both optimized (red) and trilinear basis functions (green), tested on 2 different examples of material composition using a neo-Hookean model. Note that a large deformation of an ABAB composite structure (left) is not well captured by our basis functions, while an ABBA composite structure (right) is free of such issue.

et al. 2017; Kharevych et al. 2009]: phase errors will accumulate over time (see Fig. 15). This is to be expected since we do not coarsen the inertia and the associated basis functions for the velocity field, and our shape functions are optimized in the rest state: these two facts can lead to an inaccurate capturing of the natural frequency of the dynamical system. As reported in previous works, this rarely affects the visual impact of simulation. Fabrication methods may require further improvements in this direction [Chen et al. 2017].

4.6 Performance and timing

The offline computation of our numerical coarsening approach requires solving for $d(d+1)/2$ global harmonic displacements and optimizing the coarse element shape functions. The cost of the first step is mainly related to the size of the fine mesh and each harmonic displacement can be computed in parallel. The computational time required for the shape function optimization for a given coarse element is strongly influenced by the ratio between coarse and fine




Fig. 15. **Coarsened dynamics.** Our method (top right) also captures the dynamic behavior of a neo-Hookean composite material (top left)). The curves (bottom) show the vertical displacement of the light-blue node in time. Comparing with simple FEM coarsening (green), our result (blue) matches the one of the fine mesh (red) much more closely.




Fig. 16. **Falling heart.** The heart model with an empty chamber, whose fine discretization has 89400 cells and 161060 nodes, is falling on the ground. For a one-level coarsened model with 11175 cells and 28942 nodes, the simulation using our optimized basis functions (middle) looks nearly identical (left), while trilinear basis functions (right) lead to obvious inaccuracies.

grid elements; for instance, a matrix-valued shape function has 1944 (resp., 9000) DOFs for a coarse element containing 2^3 (resp., 4^3) fine elements respectively. We provide timings for a single coarse element’s optimization in Section 4.6 for various examples shown in this paper. Note finally that all coarse elements can be processed in parallel, and the computational complexity for basis optimization only depends on the number of coarsening levels.

At runtime, positional constraints (used to anchor selected nodes to a given position) are implemented as penalty, and Newton’s method with trust region is used to solve the resulting unconstrained optimization problem of Eq. (1). We terminate the Newton iterations when the infinite norm of the gradient is less than $1e^{-6}$ or the number of iterations exceeds 20. Directly solving the problem on a fine mesh using classical FEM is obviously slower than our coarsened treatment by at least a factor proportional to the ratio between the coarse and the fine grid sizes. When each coarse element contains 8 fine elements (that is, each dimension was coarsened by a factor two), solving each linear system involved in a Newton iteration takes 20 to 50 times less on the coarse mesh than on the fine mesh. This acceleration can reach over 1000 times when each dimension is coarsened by a factor 4. Moreover, convergence is typically faster on coarse meshes (a well-known property exploited in multigrid methods); our experiments confirm this fact too: as shown in Fig. 17, the solver converges faster on the coarse mesh. The acceleration factors due to faster linear solves are thus very predictive lower bounds for the acceleration in wall-clock time of our coarsening approach. Overall, a single-level coarse simulation runs approximately

Table 1. **Statistics.** Timing (in seconds) of offline computations on a Xeon(R) E5-2630 CPU. Complexity (in elements and vertices) of the fine and coarse meshes are shown for each model. Columns denoted h_{ab} and n_{ij} indicate the time needed to solve *one* global harmonic displacement on the fine mesh and the shape functions’ optimization for *one* coarse element respectively.

Model	$\#\Omega_e^h$	$\#\text{node in } \Omega^h$	$\#\Omega_e^H$	h_{ab}	n_{ij}
cube	4096	4913	64	2.15	7.5
aircraft	21328	26801	2666	29.82	0.25
bridge	5632	7565	704	2.79	0.25
beam	2048	2673	256	0.82	0.25
hand	26176	31337	3272	35.01	0.25

60 times faster than on the fine mesh, without significant visible error; a two-level simulation typically runs over 1000 times faster than its fine simulation.

Finally, the dynamic simulation in Fig. 15 is performed through a variational implicit integrator [Martin et al. 2011], and achieves similar speedups to the static case. Because of inertia, the deformation between two frames is relatively small, so the Newton solver usually takes only 10 iterations on the fine mesh and 7 iterations on the coarsened mesh.

5 LIMITATIONS AND FUTURE WORK

Our numerical coarsening through the optimization of shape functions brings significant improvements upon previous methods focusing on capturing complex deformation at low computational cost. In particular, combining the use of matrix-valued shape functions and their element-wise corotational frames offers the ability to explore the careful balance between homogenization of each element’s behavior and inter-element induced stiffness during simulation.

While our approach can capture composite materials made of non-linear constitutive laws, we believe that a number of variants

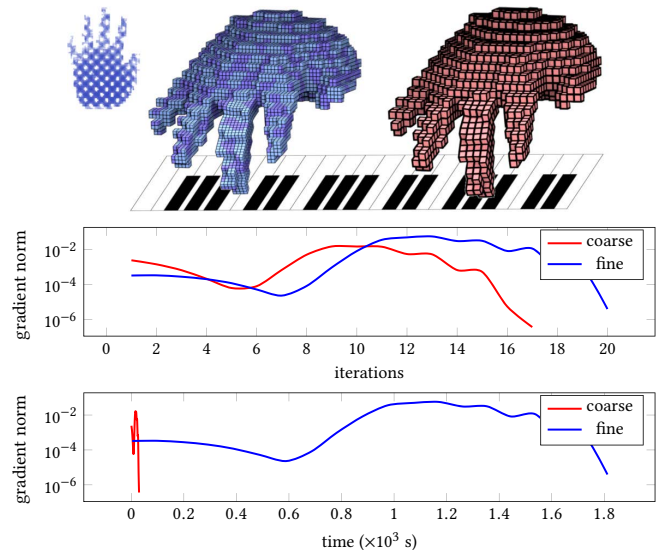


Fig. 17. **Hand playing piano.** For a coarsened hand model (right) whose fingers dynamically bend under gravity, convergence of the solver based on the (infinite norm of the) gradient (bottom) happens much earlier than for the fine model (left), and takes about 1/60-th of the computational time.

and extensions would be interesting to explore, and could further alleviate the current limitations. First, the use of analytical shape functions may be sufficient for simple materials, removing the need for a large amount of DOFs and potentially allowing realtime use in gaming. The spatial discretization of shape functions may at least be adapted based on the local non-linearity of the material to reduce the DOFs that need to be optimized without losing too much accuracy in the process. Second, better cubature schemes and solvers may improve solver efficiency further. Third, providing a robust approach to numerical coarsening of dynamics using inhomogeneous mass matrices would be an important extension that we have not attempted to address. We believe that space-time coarsening may be necessary to achieve accurate results. Fourth, applying our ideas to acoustics or other physical modeling tasks would also be valuable.

ACKNOWLEDGMENTS

We would like to thank all the anonymous reviewers for their valuable comments and suggestions. This work was partially supported by the NSFC (Nos. 61522209, 61210007) and Fundamental Research Funds for the Central Universities (2018FZA5011). MD gratefully acknowledges the INRIA International Chair program; he also thanks Max Budninskiy, Houman Owahdi, and Guoyu Lin for feedback, and Zhejiang University for hosting him during the final revisions of this paper.




Fig. 18. **Stress distribution.** The trace of the (first Piola-Kirchhoff) stress tensor exhibits a similar distribution for fine vs. coarsened meshes. On the plane example of Fig. 1, we observe the expected large stretches (positive trace on wing top, left) and compressions (negative trace on wing bottom, right) over the parts of the plane containing the stiffer material.

REFERENCES

- Jernej Barbič and Doug L. James. 2005. Real-Time Subspace Integration for St. Venant-Kirchhoff Deformable Models. *ACM Transactions on Graphics (TOG)* 24, 3 (July 2005), 982–990.
- Alain Bensoussan, Jacques-Louis Lions, and George Papanicolaou. 1978. *Asymptotic analysis for periodic structures*. Vol. 5. North-Holland Publishing Company.
- Desai Chen, David IW Levin, Wojciech Matusik, and Danny M Kaufman. 2017. Dynamics-aware numerical coarsening for fabrication design. *ACM Transactions on Graphics (TOG)* 36, 4, Article 84 (2017), 15 pages.
- Desai Chen, David IW Levin, Shinjiro Sueda, and Wojciech Matusik. 2015. Data-driven finite elements for geometry and material design. *ACM Transactions on Graphics (TOG)* 34, 4, Article 74 (2015), 10 pages.
- Bernardo Cockburn, George E. Karniadakis, and Chi-Wang Shu. 2011. *Discontinuous Galerkin Methods: Theory, Computation and Applications*. Springer.
- Tim Davis. 2017. Suitesparse: a suite of sparse matrix software. Found at <http://faculty.cse.tamu.edu/davis/suitesparse.html>.
- 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*. ACM, New York, NY, USA, 31–36.
- Yalchin Efendiev, Juan Galvis, and Xiao-Hui Wu. 2011. Multiscale finite element methods for high-contrast problems using local spectral basis functions. *J. Comput. Phys.* 230, 4 (2011), 937–955.
- Elliot English and Robert Bridson. 2008. Animating developable surfaces using non-conforming elements. *ACM Transactions on Graphics (TOG)* 27, 3, Article 66 (2008), 5 pages.
- François Faure, Benjamin Gilles, Guillaume Bousquet, and Dinesh K. Pai. 2011. Sparse meshless models of complex deformable solids. *ACM transactions on graphics (TOG)* 30, 4, Article 73 (2011), 10 pages.
- Benjamin Gilles, Guillaume Bousquet, Francois Faure, and Dinesh K. Pai. 2011. Frame-based Elastic Models. *ACM Transactions on Graphics (TOG)* 30, 2, Article 15 (2011), 12 pages.
- Eitan Grinspun, Petr Krysl, and Peter Schröder. 2002. CHARMS: A Simple Framework for Adaptive Simulation. *ACM Transactions on Graphics* 21, 3 (July 2002), 281–290.
- Geoffrey Irving, Joseph Teran, and Ronald Fedkiw. 2006. Tetrahedral and hexahedral invertible finite elements. *Graphical Models* 68, 2 (2006), 66–89.
- A. Johnen, J.-F. Remacle, and C. Geuzaine. 2013. Geometrical validity of curvilinear finite elements. *J. Comp. Phys.* 233, Supplement C (2013), 359–372.
- Lily Kharevych, Patrick Mullen, Houman Owahdi, and Mathieu Desbrun. 2009. Numerical coarsening of inhomogeneous elastic materials. *ACM Transactions on Graphics (TOG)* 28, 3, Article 51 (2009), 8 pages.
- 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.
- 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 Transactions on Graphics (TOG)* 33, 4, Article 108 (2014), 10 pages.
- Sebastian Martin, Peter Kaufmann, Mario Botsch, Martin Wicke, and Markus Gross. 2008. Polyhedral finite elements using harmonic basis functions. *Computer Graphics Forum* 27, 5 (2008), 1521–1529.
- Sebastian Martin, Bernhard Thomaszewski, Eitan Grinspun, and Markus Gross. 2011. Example-based elastic materials. *ACM Transactions on Graphics (TOG)* 30, 4, Article 72 (2011), 8 pages.
- Aleka McAdams, Yongning Zhu, Andrew Selle, Mark Empey, Rasmus Tamstorf, Joseph Teran, and Efthychios Sifakis. 2011. Efficient elasticity for character skinning with contact and collisions. *ACM Transactions on Graphics (TOG)* 30, 4 (2011), 37.
- Matthias Müller and Markus Gross. 2004. Interactive Virtual Materials. In *Proceedings of Graphics Interface 2004 (GI '04)*. Canadian Human-Computer Communications Society, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 239–246.
- Rahul Narain, Armin Samii, and James F. O'Brien. 2012. Adaptive anisotropic remeshing for cloth simulation. *ACM Transactions on Graphics (TOG)* 31, 6, Article 152 (2012), 10 pages.
- Matthieu Nesme, Paul G Kry, Lenka Jeřábková, and François Faure. 2009. Preserving topology and elasticity for embedded deformable models. *ACM Transactions on Graphics (TOG)* 28, 3, Article 52 (2009), 9 pages.
- Houman Owahdi, Lei Zhang, and Leonid Berlyand. 2014. Polyharmonic homogenization, rough polyharmonic splines and sparse super-localization. *ESAIM: Mathematical Modelling and Numerical Analysis* 48, 2 (2014), 517–552.
- Julian Panetta, Qingnan Zhou, Luigi Malomo, Nico Pietroni, Paolo Cignoni, and Denis Zorin. 2015. Elastic Textures for Additive Fabrication. *ACM Transactions on Graphics (TOG)* 34, 4, Article 135 (2015), 12 pages.
- Alex Pentland and John Williams. 1989. Good Vibrations: Modal Dynamics for Graphics and Animation. *SIGGRAPH Comput. Graph.* 23, 3 (1989), 207–214.
- Gangan Prathap. 1993. *The Finite Element Method in Structural Mechanics*. Solid Mechanics and Its Applications, Vol. 24. Springer.
- Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. 2015. Microstructures to control elasticity in 3D printing. *ACM Transactions on Graphics (TOG)* 34, 4, Article 136 (2015), 13 pages.
- Jonathan R. Shewchuk. 2002. What is a Good Linear Element? Interpolation, Conditioning, and Quality Measures. In *Int. Meshing Roundtable*. 115–126.
- 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).
- Rosell Torres, Alejandro Rodríguez, José M. Espadero, and Miguel A. Otaduy. 2016. High-resolution interaction with corotational coarsening models. *ACM Transactions on Graphics (TOG)* 35, 6, Article 211 (2016), 11 pages.