DOI: 10.1111/cgf.14622
ACM SIGGRAPH / Eurographics Symposium on Computer Animation 2022
D. L. Michels and S. Pirk
(Guest Editors)
Volume 41 (2022), Number 8

Voronoi Filters for Simulation Enrichment

Juan J. Casafranca[1,2] and Miguel A. Otaduy[1]

[1]Universidad Rey Juan Carlos, Madrid, Spain
[2]SEDDI, Madrid, Spain

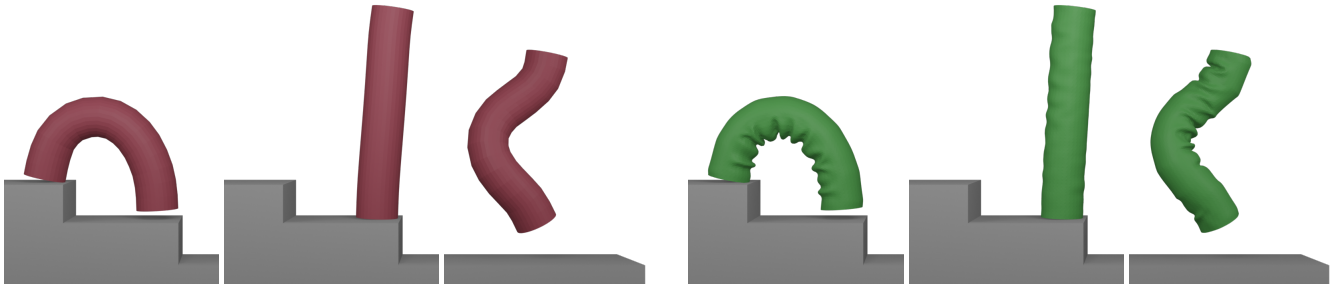

Figure 1: *We start by animating a coarse cylinder (left) using a skinning rig. Then we enhance the animation by enriching the surface with a fine thin-shell model (right). Voronoi filters provide an easy and efficient way of enriching the input animation.*

Abstract

The simulation of complex deformation problems often requires enrichment techniques that introduce local high-resolution detail on a generally coarse discretization. The use cases include spatial or temporal refinement of the discretization, the simulation of composite materials with phenomena occurring at different scales, or even codimensional simulation. We present an efficient simulation enrichment method for both local refinement of the discretization and codimensional effects. We dub our method Voronoi filters, as it combines two key computational elements. One is the use of kinematic filters to constrain coarse and fine deformations, and thus provide enrichment functions that are complementary to the coarse deformation. The other one is the use of a centroidal Voronoi discretization for the design of the enrichment functions, which adds high-resolution detail in a compact manner while preserving the rigid modes of coarse deformation. We demonstrate our method on simulation examples of composite materials, hybrid triangle-based and yarn-level simulation of cloth, or enrichment of flesh simulation with high-resolution detail.

CCS Concepts

• ***Computing methodologies*** → ***Physical simulation;***

1. Introduction

In their simplest form, simulations of deformable objects use a single fixed discretization and a single material model (possibly heterogeneous) per object. Simulations become more complex when discretizations must be spatially or temporally adapted [MWN*17, NSO12], or when multiple material models are combined, e.g., to represent composite materials [Gib10], or combinations of phenomena best defined at different dimensionality [MKB*10,LKJ21]. In this work, we use the term *enrichment* to denote two types of simulations. One is a refinement of the discretization. Another one is the simulation of composite models where a substrate model is combined with another possibly local model, which describes a phenomenon occurring at higher resolution. In both cases, the coarser discretization of the substrate is combined with a finer discretization to describe the high-resolution phenomena.

We present a simulation enrichment method to handle in a computationally efficient manner both discretization refinement and composite models. Our formulation builds on the method of *kinematic filters* [BW98,AB03,MGL*15,CCR*20]. To enrich a coarse displacement, we start by simply superposing a fine displacement, but then we constrain both coarse and fine displacements to be orthogonal and thus remove redundancies. This constraint is efficiently imposed using kinematic filters. In Section 3, we describe the general formulation of enrichment using kinematic filters, and

we compare it to a regular constrained-based dynamics formulation. To the best of our knowledge, previous works cover only bits and pieces or special cases of the kinematic-filter formulation. Our general exposition allows us to list design objectives for efficient and robust enrichment.

Based on these objectives, in Section 4, we introduce *centroidal Voronoi enrichment*. Our enrichment method enjoys non-overlapping coarse basis functions, which enable extremely fast implementation of the kinematic filter. However, naïve non-overlapping basis functions do not satisfy rigid invariance, i.e., it is not possible to achieve rigid motion on both the coarse and fine discretizations at the same time, and rigid invariance is key to ensure that no spurious forces are injected between coarse and fine discretizations. We show that a simple and elegant modification to the coarse mesh allows our enrichment to satisfy rigid invariance.

We dub the combination of centroidal Voronoi enrichment and kinematic filters as *Voronoi filters*. In the paper, we showcase Voronoi filters on a variety of applications. The examples include composite materials with codimensional embedded objects, mixed simulation of triangle-based and yarn-level cloth, and enhancement of coarse simulations with high-resolution surface detail, as in Fig. 1. Voronoi filters are very simple to implement, and the enriched simulation bears practically no overhead with respect to unconstrained simulation.

2. Related Work

In computer graphics, simulation enrichment is typically connected to the design of adaptive meshing or mesh refinement methods [WDGT01, DDCB01, MWN*17], formally known as *h*-refinement. This approach has obtained success for triangle-based simulation of cloth [NSO12], where the construction of a consistent adaptive mesh can be solved through local operations. Mesh refinement for volumetric meshes requires however more complex solutions, such as higher-order basis functions, discussed next.

In contrast to mesh refinement, an alternative to obtain higher accuracy in simulations of deformable objects is to use basis functions of higher order, formally known as *p*-refinement. This approach has recently gained popularity to enforce accuracy bounds on meshes of bad quality [SHD*18], for isogeometric analysis where simulation meshes conform to the curved boundaries used in CAD applications [SDG*19], or to approximate high-resolution surfaces with coarse simulation meshes with curved boundaries [LLK*20, JZH*21].

A combination of both types of refinement considers the addition of basis functions of progressively higher resolution, known as *hp*-refinement. Such basis functions may be designed through subdivision schemes [GKS02], ensuring that the refinement bases are orthogonal to the underlying bases, i.e., fine basis functions capture only displacements that are not captured by coarse basis functions. A related approach has recently been followed for the design of material-adapted refinable basis functions [CBO*19].

Mesh and/or basis refinement methods all assume that the simulation models are defined on a simulation domain of a given dimensionality. However, they do not consider codimensional objects or composite materials with simulation models best described on different dimensionalities. In our work, we investigate a refinement method that addresses both the local addition of detail as well as the combination of models of different dimensionality.

An alternative to refinement is to combine discretizations of different resolutions (or dimensionality). If no special care is taken, these discretizations produce a redundant kinematic space. Therefore, constraints must be added to remove the redundancy. Bergou et al. [BMWG07] proposed a constrained optimization to solve such constraints, and produce a high-resolution simulation that matches the overall deformation of a coarse simulation.

Malgat et al. [MGL*15] introduced kinematic filters to solve the redundancy constraints efficiently, and they demonstrated both refined simulations and embedded codimensional effects. Casafranca et al. [CCR*20] used kinematic filters to combine triangle-based and yarn-level models in cloth simulation. These works focused on the development of the kinematic filter approach and formulating and solving dynamics with such filters, but they did not pay particular attention to the design of the filters. Our work also builds on the method of kinematic filters, but we introduce a novel filter that achieves high computational efficiency thanks to a diagonal kernel. Other works have also used ideas similar to kinematic filters to transfer data between representations, e.g, between grids and particles [JSS*15], or between grids and rigid modes [FLLP13]. Some recent works constrain motion to the null-space of predefined reduced spaces; some use kinematic filters to implement the constraint efficiently [SYS*21, RCPO22], while others could replace Lagrange multiplier formulations with kinematic filters for higher performance [ZBLJ20].

Several works have paid attention to the combination of models formulated on different dimensionality, i.e., volumetric solids, shells, and rods. The standard approach is to conform the high-dimensional mesh to the embedded low-dimensional elements [HZG*18], but we avoid the need to mesh the high-dimensional domain at a high resolution imposed by the embedded low-dimensional elements. Martin et al. [MKB*10] designed integration rules that handle in a unified manner deformation models of any dimensionality. Macklin et al. [MMCK14] supported multi-physics effects using a particle-based discretization and constraint-based models as unifying framework. Recently, Chang et al. [CDGB19] have presented a general energy formulation to handle the connection of objects with different dimensionality. These methods require either conforming discretizations for the objects with different dimensionality, or some type of explicit constraint between them. When the constraints are hard, they require constrained optimization solvers. Our method does not require conforming meshes for codimensional objects, and handles constraints efficiently thanks to a kinematic filter.

3. Enrichment with Kinematic Filters

In this section, we present the fundamentals of kinematic filters as a way to implement simulation enrichment. We first present general enrichment with non-redundancy constraints. Then we derive dynamics equations, and we compare the solution approach with regular constrained dynamics vs. the use of kinematic filters. We

conclude with a discussion of the computational cost of kinematic filters and design desiderata that motivate our formulation.

3.1. Enriched Kinematics

We start from a coarse discretization that defines a coarse displacement field. This coarse displacement field may describe the deformation of a substrate in the case of composite models, or a low-resolution version of the full deformation in the case of refinement problems. We wish to enrich this displacement field with displacements computed on a fine discretization. The enrichments may describe the offset deformation of an embedded local material in the case of composite models, or the high-resolution detail in the case of refinement problems. Without loss of generality, we will analyze the displacements at fine node locations; these can be then interpolated at arbitrary points.

Let us define the deformed and undeformed positions of fine nodes as $\{\mathbf{x}_i\}$ and $\{\bar{\mathbf{x}}_i\}$ respectively. Similarly, let us define coarse node displacements $\{\mathbf{u}^c_j\}$, fine node enrichments $\{\mathbf{u}^f_i\}$, and full fine node displacements $\{\mathbf{u}_i\}$.

We define the displacement of a fine node as the sum of a coarse displacement and a fine node enrichment. The coarse displacement field may be constructed from coarse node displacements using coarse basis functions $\{\mathbf{B}_j(\bar{\mathbf{x}})\}$. Then, the full displacement of a fine node can be formally expressed as:

$$\mathbf{u}_i = \sum_j \mathbf{B}_{ij}\,\mathbf{u}^c_j + \mathbf{u}^f_i, \tag{1}$$

where $\mathbf{B}_{ij} = \mathbf{B}_j(\bar{\mathbf{x}}_i)$ is the evaluation of a coarse basis function at a fine node location.

We rewrite (1) in matrix-vector form, with concatenated vectors of coarse displacements $\mathbf{u}^c$, enrichments $\mathbf{u}^f$, fine node displacements $\mathbf{u}$, and a matrix of coarse basis function values $\mathbf{B}$:

$$\mathbf{u} = \mathbf{B}\,\mathbf{u}^c + \mathbf{u}^f. \tag{2}$$

The enrichment function (2) produces a redundant kinematic representation, as infinite combinations of coarse displacements and fine enrichments produce the same full displacements. We want the enrichment to be complementary to the coarse displacement, and we achieve this by constraining the enrichment to be orthogonal to the coarse basis function:

$$\mathbf{B}^T\,\mathbf{u}^f = 0. \tag{3}$$

Since the constraint matrix $\mathbf{B}^T$ is constant, an alternative to enforcing constraints on every simulation frame would be to define enrichments using as basis the orthogonal complement of $\mathbf{B}$, e.g., through QR decomposition. However, this basis is dense, and it would lead to dense energy Hessians. Kinematic filters achieve the same goal without ever constructing a dense basis.

3.2. Enriched Dynamics with Constraints

In a general case, we consider inertial and elastic effects defined at both coarse and fine resolution, e.g., for our codimensional simulation examples. Then, we define kinetic energy $\frac{1}{2}\,\dot{\mathbf{u}}^c\,\mathbf{M}^c\,\dot{\mathbf{u}}^c + \frac{1}{2}\,\dot{\mathbf{u}}\,\mathbf{M}^f\,\dot{\mathbf{u}}$, with coarse and fine mass matrices $\mathbf{M}^c$ and $\mathbf{M}^f$ respectively, and elastic energy $U^c(\mathbf{u}^c) + U^f(\mathbf{u})$. Note that fine-mesh kinetic energy depends on full fine node velocities $\dot{\mathbf{u}}$, not just enrichment velocities $\dot{\mathbf{u}}^f$.

Without loss of generality, we formulate dynamics using the optimization formulation of backward Euler integration [MTGG11, GSS*15]. We denote as $\mathbf{u}^{c*}$ and $\mathbf{u}^*$ the explicit update of coarse and fine displacements, and h the time step. The update of coarse displacements and fine enrichments is computed as the constrained optimization

$$\begin{aligned}(\mathbf{u}^c,\mathbf{u}^f) = \arg\min\ & \frac{1}{2h^2}\,(\mathbf{u}^c-\mathbf{u}^{c*})^T\,\mathbf{M}^c\,(\mathbf{u}^c-\mathbf{u}^{c*}) + U^c(\mathbf{u}^c)\\ & + \frac{1}{2h^2}\,(\mathbf{u}-\mathbf{u}^*)^T\,\mathbf{M}^f\,(\mathbf{u}-\mathbf{u}^*) + U^f(\mathbf{u}),\\ \text{s.t. } & \mathbf{B}^T\,\mathbf{u}^f = 0,\end{aligned} \tag{4}$$

with fine displacements $\mathbf{u}$ defined as in (2).

This type of constrained dynamics problem is typically solved using Lagrange multipliers, e.g., as in [BMWG07]. However, constraints such as (3) can also be enforced using a projection operation [BW98], which leads to the definition of kinematic filters.

3.3. Solver with Kinematic Filter

Let us define filtered enrichments $\tilde{\mathbf{u}}^f$ as those closest to tentative enrichments $\mathbf{u}^f$, but which satisfy the constraint (3). Formally, this can be expressed as $\tilde{\mathbf{u}}^f = \arg\min \|\tilde{\mathbf{u}}^f - \mathbf{u}^f\|^2$, s.t. $\mathbf{B}^T\,\tilde{\mathbf{u}}^f = 0$. The result is given by filtered enrichments $\tilde{\mathbf{u}}^f = \mathbf{F}\,\mathbf{u}^f$, with a kinematic filter $\mathbf{F} = \mathbf{I} - \mathbf{B}\left(\mathbf{B}^T\,\mathbf{B}\right)^{-1}\mathbf{B}^T$.

The enriched kinematics (2) are reformulated as:

$$\mathbf{u} = \mathbf{B}\,\mathbf{u}^c + \mathbf{F}\,\mathbf{u}^f. \tag{5}$$

Note that the filter projects the enrichment to the nullspace of the coarse basis functions, and this effectively removes the redundancy in the kinematic representation.

With the filtered kinematics, the constrained dynamics (4) can be reformulated as an unconstrained optimization:

$$\begin{aligned}(\mathbf{u}^c,\mathbf{u}^f) = \arg\min\ & \frac{1}{2h^2}\,(\mathbf{u}^c-\mathbf{u}^{c*})^T\,\mathbf{M}^c\,(\mathbf{u}^c-\mathbf{u}^{c*}) + U^c(\mathbf{u}^c)\\ & + \frac{1}{2h^2}\,(\mathbf{u}-\mathbf{u}^*)^T\,\mathbf{M}^f\,(\mathbf{u}-\mathbf{u}^*) + U^f(\mathbf{u}),\end{aligned} \tag{6}$$

with fine displacements $\mathbf{u}$ now defined as in (5).

Computing the gradient of the objective (6) yields the backward Euler update:

$$\begin{aligned}\mathbf{M}\begin{pmatrix}\Delta\mathbf{u}^c\\ \Delta\mathbf{u}^f\end{pmatrix} &= -h^2\begin{pmatrix}\nabla U^c + \mathbf{B}^T\,\nabla U^f\\ \mathbf{F}^T\,\nabla U^f\end{pmatrix},\\ \text{with } \mathbf{M} &= \begin{pmatrix}\mathbf{M}^c + \mathbf{B}^T\,\mathbf{M}^f\,\mathbf{B} & \mathbf{B}^T\,\mathbf{M}^f\,\mathbf{F}\\ \mathbf{F}^T\,\mathbf{M}^f\,\mathbf{B} & \mathbf{F}^T\,\mathbf{M}^f\,\mathbf{F}\end{pmatrix}.\end{aligned} \tag{7}$$

Applying Newton-type methods yields a Hessian of the objective that is not full rank, due to the kinematic filter. However, the resulting linear system can be solved robustly and efficiently with the modified preconditioned conjugate gradient (MPCG)

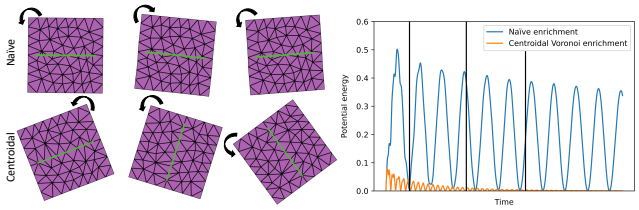

Figure 2: *In this test, we spin a deformable square coupled to a rod. On top, we use naïve enrichment as described in Section 4.1 to define the kinematics of the rod. When the rod spins with constant angular velocity, it produces a deformation on the square. Due to ghost forces, the square oscillates and cannot spin. In the bottom, we use centroidal Voronoi enrichment, and the square spins correctly. At the beginning it shows some deformation energy, produced by the initial acceleration. The snapshots are taken at the instants marked by the black lines in the plot.*

method [AB03]. On each iteration of conjugate gradient, the enrichments are filtered with the filter $\mathbf{F}$.

In the examples shown in the paper, we also add damping to the dynamics. We do this using the first-order-accurate extension of backward-Euler optimization proposed by Brown et al. [BOFN18], and with damping models designed using the dissipation potentials proposed by Sánchez-Banderas and Otaduy [SBO18].

3.4. Discussion and Desiderata

The cost of solving enriched dynamics with kinematic filters is determined by the number of iterations of MPCG and the cost per iteration. Each iteration requires a matrix-vector multiplication with the Hessian of (6) *and* with the filter. This multiplication with the filter requires in turn a reduction of fine node enrichments, followed by a solution of a coarse linear problem with the matrix $\mathbf{B}^T\mathbf{B}$, and finally an expansion back to fine displacements. In previous kinematic filters for deformation gradients [MGL*15] or barycentric interpolation [CCR*20], solving the linear system with $\mathbf{B}^T\mathbf{B}$ becomes a bottleneck as the resolution of the coarse discretization grows.

In designing enrichment functions, we wish to fulfill two goals: (1) The cost of the resulting kinematic filter $\mathbf{F}$ should be minimized. (2) The coarse basis $\mathbf{B}$ should preserve rigid motion, to ensure that spurious deformations are not injected between the coarse and fine discretizations.

4. Centroidal Voronoi Enrichment

In this section, we describe our simulation enrichment method. We start by formulating Voronoi enrichment as a method to maximize the efficiency of kinematic filters. Then we impose rigid invariance, which leads to centroidal Voronoi enrichment. We conclude the section with a discussion of implementation details of the discretization.

4.1. Voronoi Enrichment

We can maximize the efficiency of kinematic filters by choosing simple zeroth-order coarse basis functions. Specifically, we set the support of each coarse node $\mathbf{x}_j$ as its Voronoi region; i.e., the basis function of $\mathbf{x}_j$ is identity at all fine nodes $\mathbf{x}_i$ within its Voronoi region:

$$\mathbf{B}_{ij} = \begin{cases} \mathbf{I}, & \text{if } j = \arg\min_j |\bar{\mathbf{x}}_i - \bar{\mathbf{x}}_j| \\ \mathbf{0}, & \text{otherwise} \end{cases} \tag{8}$$

The use of zeroth-order functions introduces discontinuities in how the coarse displacement field is interpolated to the fine nodes, but the fine displacement field does not suffer discontinuties, thanks to the enrichment degrees of freedom.

With these basis functions, the product with the resulting kinematic filter has a negligible cost. $\mathbf{B}^T\mathbf{B}$ is a diagonal matrix, with entries N_j (the number of fine nodes in each coarse node's Voronoi region), hence it is trivially invertible. In practice the global Voronoi filter is composed of compact separable filters (i.e., small and uncoupled), one for the Voronoi region of each coarse node. Given tentative fine enrichments $\mathbf{u}^f$, filtering amounts to computing the average enrichment per Voronoi region, and subtracting these averages from the enrichments. As we discuss in the results section, the cost of this operation is negligible in practice.

We can also analyze the effect of Voronoi filters on the full fine node displacements $\{\mathbf{u}_i\}$. Applying the constraint (3) with the basis (8), we obtain $\sum_i \mathbf{u}_i^f = 0$, i.e., the average enrichment is zero. Substituting this into the expression of the fine node displacement (1), and averaging over all fine nodes, we obtain:

$$\mathbf{u}_j^c = \frac{1}{N_j}\sum_i \mathbf{u}_i. \tag{9}$$

In other words, with Voronoi filters, the displacement of a coarse node is the average displacement of the fine nodes within its Voronoi region.

4.2. Rigid Invariance

To blend elasticity models defined at different scales or on codimensional domains, we also wish to preserve rigid modes on both the fine and coarse discretization. Unfortunately, with a naïve implementation of enrichment, based on zeroth-order coarse basis functions, a rigid motion of the fine nodes could inject non-rigid motion on coarse nodes through the filter (see the averaging effect in (9), or viceversa. This would produce elastic energy on the coarse model, and therefore ghost forces on the simulation. An example is discussed in Fig. 2.

Next, we derive a sufficient and necessary condition on the undeformed positions of coarse nodes for the preservation of rigid modes. In a nutshell, each coarse node must be located at the centroid of its Voronoi region. Therefore, the coarse discretization defines a discrete centroidal Voronoi tessellation [DFG99] of the fine discretization.

Let us define a rigid transformation by a rotation matrix $\mathbf{R}$ and a translation vector $\mathbf{t}$. If both coarse and fine nodes are transformed

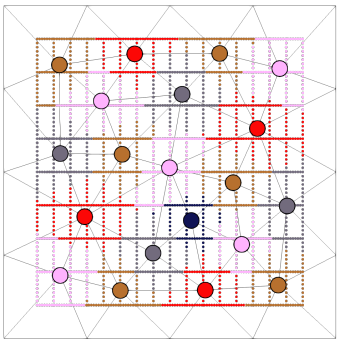

Figure 3: *This image shows the construction of a discretization with Voronoi enrichment. A coarse triangle mesh is enriched with a grid of rods, discretized at higher resolution. We color each coarse node and the fine nodes in its Voronoi region with the same color. To guarantee rigid-motion invariance, the coarse mesh constitutes a centroidal Voronoi diagram.*

with the same rigid transformation, we can obtain their displacements as:

$$\mathbf{x}_j = \bar{\mathbf{x}}_j + \mathbf{u}_j^c = \mathbf{R}\bar{\mathbf{x}}_j + \mathbf{t} \quad \Rightarrow \quad \mathbf{u}_j^c = (\mathbf{R}-\mathbf{I})\bar{\mathbf{x}}_j + \mathbf{t} \tag{10}$$

$$\mathbf{x}_i = \bar{\mathbf{x}}_i + \mathbf{u}_i = \mathbf{R}\bar{\mathbf{x}}_i + \mathbf{t} \quad \Rightarrow \quad \mathbf{u}_i = (\mathbf{R}-\mathbf{I})\bar{\mathbf{x}}_i + \mathbf{t}. \tag{11}$$

We know that, due to the Voronoi filter, the displacement of the coarse node is the average displacement of fine nodes (9). Applying this to (10) and (11), we obtain:

$$(\mathbf{R}-\mathbf{I})\bar{\mathbf{x}}_j + \mathbf{t} = \frac{1}{N_j}\sum_i \left((\mathbf{R}-\mathbf{I})\bar{\mathbf{x}}_i + \mathbf{t}\right). \tag{12}$$

From this, we derive a trivial condition on translations, which is satisfied by construction, $\mathbf{t} = \frac{1}{N_j}\sum_i \mathbf{t}$. Most importantly, we also derive a condition on rotations, which leads to the constraint on undeformed positions of coarse nodes:

$$(\mathbf{R}-\mathbf{I})\bar{\mathbf{x}}_j = (\mathbf{R}-\mathbf{I})\frac{1}{N_j}\sum_i \bar{\mathbf{x}}_i \quad \Rightarrow \quad \bar{\mathbf{x}}_j = \frac{1}{N_j}\sum_i \bar{\mathbf{x}}_i. \tag{13}$$

As anticipated, to preserve rigid motion of coarse and fine discretizations, coarse nodes must lie at the centroid of their fine nodes. In other words, rigid invariance is satisfied iff the undeformed coarse discretization defines a discrete centroidal Voronoi tessellation of the undeformed fine discretization. The test in Fig. 2 works perfectly using centroidal Voronoi enrichment and Voronoi filters.

4.3. Implementation of the Discretization

The implementation of the discrete centroidal Voronoi discretization is simple, and it is outlined in Fig. 3. It involves remeshing the interior of the coarse mesh as a discrete centroidal Voronoi diagram of the fine mesh. Note that we do not remesh the boundary of the coarse mesh, as we want to preserve this boundary.

Coarse nodes	544	2112	8320	33024
% of MPCG, baryc. filter	4.8	8.0	10.7	28.4
% of MPCG, Voronoi filter	2.2	2.8	2.1	3.1

Table 1: *Comparison of cost per iteration of modified conjugate gradient [AB03] between barycentric filters [CCR*20] and Voronoi filters, for different numbers of coarse nodes.*

First, we initialize both a coarse and a fine discretization, simply by sampling nodes at two different resolutions in the region of interest. Without loss of generality, we consider the coarse discretization as a substrate that samples the complete simulation domain (i.e., the complete deformable object), while the fine discretization is possibly confined to a region in the domain.

Then, we select the coarse nodes within the fine discretization domain, excluding boundary coarse nodes, and we initialize a Voronoi diagram, i.e., for each fine node we find the closest coarse node. Subsequently, we turn it into a discrete centroidal Voronoi diagram using Lloyd's method [Llo82] on the fine nodes, as shown in Fig. 3. During Lloyd iterations, if a coarse node does not govern any fine node, i.e., $N_j = 0$, it remains fixed.

5. Experiments and Results

We have tested centroidal Voronoi enrichment on different simulation examples, which cover coupling of codimensional phenomena, hybrid simulation with different models, and simulation enhancement through mesh refinement.

5.1. Performance Analysis

We have used the example in Fig. 1 to compare the cost of Voronoi filters to barycentric filters [CCR*20]. As discussed in Section 3.4, barycentric filters introduce a small overhead when the coarse mesh is small, but the overhead grows as the nodes in the coarse mesh grow. We have confirmed this observation, by comparing simulations with barycentric and Voronoi filters on coarse meshes of different sizes. Eventually, barycentric filters can become the bottleneck of a simulation, while Voronoi filters keep a very small and constant cost.

Table 1 lists the percentage of cost of both filters per iteration of conjugate gradient solve, for different numbers of coarse nodes. As noted in the results, Voronoi filters bear a practically constant cost of just 3% of the conjugate gradient iteration.

5.2. Embedding Codimensional Phenomena

As discussed in the introduction, one of the target applications of our enrichment method is the simulation of embedded codimensional objects. In other simulation methods this requires (a) the use of constrained optimization methods, or (b) designing meshes where the discretizations of codimensional objects are conforming. Centroidal Voronoi enrichment largely simplifies the coupling of codimensional simulation models.

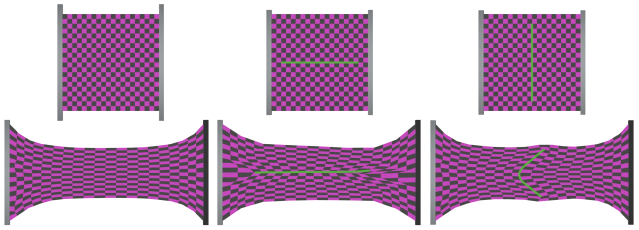

Figure 4: *The images show three deformable squares: with no embedded rods (left), with a rod embedded in the horizontal direction (middle), and with a rod embedded in the vertical direction (right). The rod is stiffer than the material of the square; therefore, it affects the overall deformation mechanics. Coupling is implemented with centroidal Voronoi enrichment.*

Fig. 4 and Fig. 5 show simple examples of embedding stiff one-dimensional rod objects into two-dimensional objects. These examples are simulated in 2D. We use a Saint Venant-Kirchhoff (StVK) energy model [OH99] for the substrate elastic energy U^c, and a discrete rod model [BWR*08] for the rod energy U^f. We simply evaluate each energy on its corresponding discretized domain and we add both energies together; the enriched kinematics produce the coupling between both models. In Fig. 4, a deformable square is pulled with different rod embeddings, showing how the embedding direction affects the overall deformation pattern. In Fig. 5, embedded rods increase the overall bending stiffness of the cantilever beam. For the horizontal rod in Fig. 4, a slight lack of symmetry is perceived in the final deformation, which is caused by the asymmetry in the underlying coarse mesh, and the stiff embedded rod brings out this asymmetry by choosing a preferred deformation.

The embedding of codimensional effects can produce interesting nonlinear behaviors on the overall simulation. Fig. 6 shows one example, where we embed a zigzagging stiff rod on a rectangular patch. When the patch is pulled, the asymmetric layout of the rod produces nonlinear wavy deformations.

Real-world composite materials are often produced by embedding codimensional objects inside two- or three-dimensional substrates. Notable examples are thin shells reinforced with carbon fiber. The proposed centroidal Voronoi enrichment offers an excellent framework to simulate such composite materials. In Fig. 7 we show an example that mimics a sheet of material reinforced with carbon fiber. We model the coarse substrate energy U^c using StVK with discrete shells bending [GHDS03], and the fine energy U^f with a grid of discrete elastic rods [BWR*08]. When the sheet is sheared, the carbon fiber reinforcement prevents the creation of folds and wrinkles, and stiffens the overall behavior as expected.

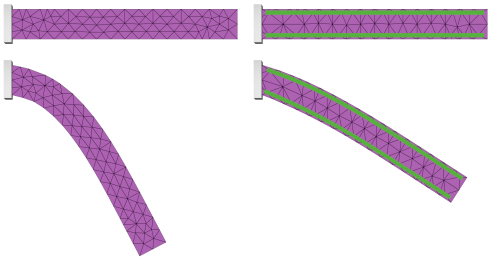

Figure 5: *The images compare a cantilever beam with two embedded rods (right) and no rods (left). The rods are stiffer than the substrate material of the beam; therefore, they increase its overall bending stiffness.*

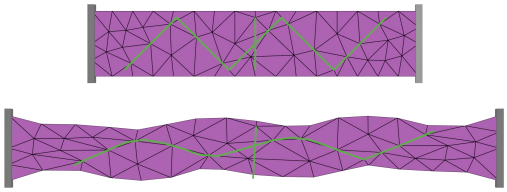

Figure 6: *This example showcases the interesting deformation effects that can be produced with embedded codimensional objects. We embed a stiff zigzagging rod inside the rectangular patch. When the patch is pulled, the rod produces a highly nonlinear wavy deformation.*

5.3. Hybrid Simulation

The simulation of hybrid or mixed models can also benefit from the proposed centroidal Voronoi enrichment. Inspired by the work of Casafranca et al. [CCR*20], we consider an example of hybrid cloth simulation where triangle-based and yarn-based models are mixed to combine their benefits: fast computation in regions with smooth deformation, and high-resolution detail in regions with wrinkles. For the coarse triangle-based model U^c, we use the same StVK material with discrete shells bending as discussed earlier, and for the fine yarn-based model U^f we use an Eulerian-on-Lagrangian approach with sliding persistent contacts [CLMMO14]. On a blending region, the triangle-based and yarn-based models are smoothly blended, and we consider these blending weights as part of the energies U^c and U^f.

Fig. 8 compares simulation results with centroidal Voronoi enrichment vs. the simulation method of Casafranca et al. The resulting deformations are practically the same, although they con-

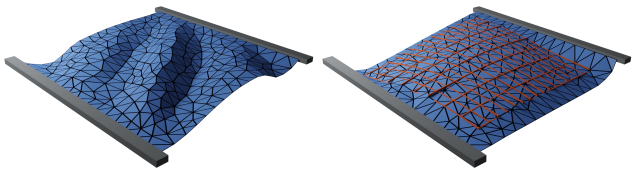

Figure 7: *Centroidal Voronoi enrichment can be used for the simulation of real-world composite materials. In this example, we demonstrate a sheet of material that mimics the construction of carbon fiber composites. A woven structure of carbon fiber, modeled as a grid of rods, is embedded in polymer resin, modeled as a deformable thin shell. The images compare a shear deformation with (right) and without (left) the carbon fiber reinforcement.*

Figure 8: *This figure shows an example borrowed from [CCR*20], which combines triangle-based and yarn-level cloth simulation. The chest area of the dress is simulated using yarns, while the rest of the dress is simulated using triangles. Both models are combined on a blending region, shown on the top left (yarns in red, triangles in green, blending in between). We have simulated this test using centroidal Voronoi enrichment, in contrast to the barycentric enrichment approach of Casafranca et al. [CCR*20]. Both methods produce almost equal results, although they converge to slightly different equilibria. However, thanks to its negligible filtering cost, centroidal Voronoi enrichment cuts computation times to about one third, as shown in the plot on the right. The image in the bottom right shows the Voronoi regions of the enrichment. The creases visible under the chest are due to seams, not a feature or limitation of our method.*

verge to slightly different equilibria. Casafranca et al. used an enrichment method based on barycentric basis functions, which produce a kinematic filter with a dense kernel. With centroidal Voronoi enrichment, the filter kernel is diagonal and trivially invertible as discussed in Section 3. As a result, the computation time of this example is cut to one third, from 21 to 7 hours. The high cost is dominated by the number of yarn-level nodes in the simulation (over 114 K) and the small time step needed (1 ms). With centroidal Voronoi enrichment, the kinematic filter takes a negligible cost of 5% per iteration of conjugate gradient. Note that in this example the hybrid region is comparatively small, as depicted on the left images in Fig. 8. The barycentric filter of Casafranca et al. would suffer an increase of computational cost with a larger hybrid region, while our centroidal Voronoi filter is barely affected.

5.4. Simulation Enhancement

One application of simulation enhancement is to execute coarse simulations until the artist is satisfied with the overall outcome, and then refine the discretization to obtain visual detail. However, the fine and coarse simulations must be constrained to ensure that the overall deformation is preserved when detail is added. Centroidal Voronoi enrichment offers a very efficient way of refining the discretization while satisfying such constraints. In the enriched dynamics equations (7), the coarse displacements are simply prescribed, and the constraints are enforced using the Voronoi filter. In this enhancement application, in the formulation of dynamics as in Section 3.2, we use coarse-mesh energies for the computation of the coarse simulation, and only fine-mesh energies (with possible blending weights) for the constrained refinement step.

In the example in Fig. 1, a cylinder is first manually animated, and then this animation is used as a constraint for a high-resolution simulation. This example is analogous to the ones demonstrated by Bergou et al. [BMWG07]. The difference is in how we constrain the coarse and fine simulations, with simple and efficient Voronoi filters. In this example, the coarse mesh has only 544 vertices; the fine mesh is simulated using a thin-shell model, and it has 8311 vertices.

In our final example, we consider a deformation problem that combines discretizations at two different resolutions, to achieve higher level of detail at localized regions of the simulation domain. Then, coarse and fine deformation energies U^c and U^f model the same problem, and they must be blended at the boundary of the fine discretization. In Fig. 9 we show a hand simulation enhanced using our method. First, we execute a coarse simulation of a finger, meshed with 1400 tetrahedra, and with animated bone transformations as boundary conditions. Then, we enrich one joint of the finger with 110K tetrahedra, and we simulate the enriched dynamics with prescribed coarse displacements. We have used the stable Neo-Hookean soft-tissue model of Smith et al. [SGK18], which produces realistic tissue bulging at the joint. In this example, the coarse simulation takes 3 minutes, and the enriched simulation takes roughly 2 hours.

6. Discussion and Future Work

This paper introduces a novel approach to couple coarse and fine discretizations. The discretizations may represent the same deformation model at different resolutions, as in adaptive simulation, or different deformation models, even of codimensional nature. The proposed centroidal Voronoi enrichment method circumvents the complexity of other methods, such as conformal mesh refinement or constraint-based optimization. It allows standard definition of the coarse and fine deformation energy models, and coupling is achieved through the enriched kinematic definition and the use of Voronoi kinematic filters during the runtime solve. The paper demonstrates the diversity of simulation examples that could benefit from our flexible enrichment method, such as the simulation of composite materials, hybrid simulations, or simulation enhancement. Furthermore, the implementation of the method is simple, and a prototype will be released.

In all the examples shown in the paper, the enrichment is fixed over time. However, the enrichment could also be updated over time to enable adaptive discretization of the simulation domain. To do this, we suggest initializing the coarse and fine discretizations once as a preprocess, and at runtime select in an adaptive man-

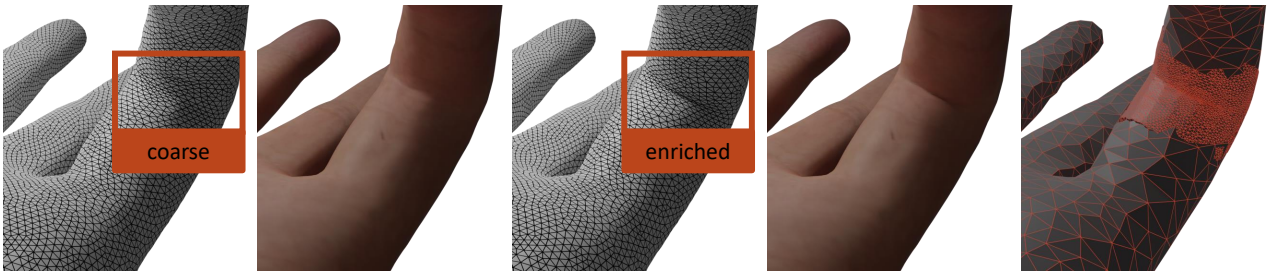

Figure 9: *We enrich a coarse tetrahedral simulation of a hand with high-resolution tetrahedra at a finger joint (right), and we compare the rendered triangle mesh before and after applying enrichment. With our Voronoi enrichment method, we can efficiently produce realistic high-resolution detail as a post-process simulation, without the need for conforming meshes or costly constrained optimizations.*

ner the coarse nodes that are enriched. Our enrichment technique is an alternative to adaptive remeshing on domains or discretization methods where remeshing is complex. Centroidal Voronoi enrichment could also be implemented on multiple discretization levels. In this case, one would start computing the centroidal Voronoi diagram on the finest discretization, and progressively move to a coarser level.

These extensions proposed above also evidence, however, the main limitation of the method. The need to compute a centroidal Voronoi diagram complicates runtime updates to the discretization. Therefore, the method is limited to simulations where the enriched discretization remains fixed over time, as shown in the examples in the paper, or where precomputation of the centroidal Voronoi diagram can be leveraged, as discussed in the paragraph above. It is unclear how to design the enrichment functions at runtime. In the same line, the method is not well suited for simulations with topology changes resulting from fracture or cutting, as they would also need a recomputation of the centroidal Voronoi diagram.

Another limitation of the method is that it is not perfectly compatible with direct solvers, e.g., based on Cholesky factorization. Note that, due to the kinematic filter, the system matrix is considerably denser than in the unfiltered version. This is not an issue with iterative solvers such as conjugate gradient, because the application of the filter at runtime is very efficient, without building the system matrix explicitly. With direct solvers, it would be convenient to design a factorization method that does not require explicitly building the system matrix. But even in that case, the resulting factorization would be considerably denser than in the unfiltered case.

Acknowledgments. We would like to thank the anonymous reviewers for their feedback. We would also like to thank the Mechanics team at SEDDI (in particular Rosa M. Sánchez-Banderas and Gabriel Cirio) and the MSLab at URJC (in particular Jesús Pérez), for support with the implementation of mechanical models. This work was funded in part by the European Research Council (ERC-2017-CoG-772738 TouchDesign).

References

[AB03] ASCHER U. M., BOXERMAN E.: On the modified conjugate gradient method in cloth simulation. *Vis. Comput. 19*, 7-8 (2003). 1, 4, 5

[BMWG07] BERGOU M., MATHUR S., WARDETZKY M., GRINSPUN E.: TRACKS: Toward Directable Thin Shells. *ACM Transactions on Graphics (SIGGRAPH) 26*, 3 (jul 2007), 50:1–50:10. 2, 3, 7

[BOFN18] BROWN G. E., OVERBY M., FOROOTANINIA Z., NARAIN R.: Accurate dissipative forces in optimization integrators. *ACM Trans. Graph. 37*, 6 (2018). 4

[BW98] BARAFF D., WITKIN A.: Large steps in cloth simulation. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques* (1998), SIGGRAPH '98, ACM, p. 43–54. 1, 3

[BWR*08] BERGOU M., WARDETZKY M., ROBINSON S., AUDOLY B., GRINSPUN E.: Discrete elastic rods. *ACM Trans. Graph. 27*, 3 (2008), 63:1–63:12. 6

[CBO*19] CHEN J., BUDNINSKIY M., OWHADI H., BAO H., HUANG J., DESBRUN M.: Material-adapted refinable basis functions for elasticity simulation. *ACM Trans. Graph. 38*, 6 (2019). 2

[CCR*20] CASAFRANCA J. J., CIRIO G., RODRÍGUEZ A., MIGUEL E., OTADUY M. A.: Mixing yarns and triangles in cloth simulation. *Computer Graphics Forum 39*, 2 (2020), 101–110. 1, 2, 4, 5, 6, 7

[CDGB19] CHANG J., DA F., GRINSPUN E., BATTY C.: A unified simplicial model for mixed-dimensional and non-manifold deformable elastic objects. *Proc. ACM Comput. Graph. Interact. Tech. 2*, 2 (2019). 2

[CLMMO14] CIRIO G., LOPEZ-MORENO J., MIRAUT D., OTADUY M. A.: Yarn-level simulation of woven cloth. *ACM Trans. Graph. 33*, 6 (Nov. 2014), 207:1–207:11. 6

[DDCB01] DEBUNNE G., DESBRUN M., CANI M.-P., BARR A. H.: Dynamic real-time deformations using space & time adaptive sampling. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2001), SIGGRAPH '01, Association for Computing Machinery, p. 31–36. 2

[DFG99] DU Q., FABER V., GUNZBURGER M.: Centroidal voronoi tessellations: Applications and algorithms. *SIAM Review 41*, 4 (1999), 637–676. 4

[FLLP13] FAN Y., LITVEN J., LEVIN D. I. W., PAI D. K.: Eulerian-on-lagrangian simulation. *ACM Trans. Graph. 32*, 3 (2013). 2

[GHDS03] GRINSPUN E., HIRANI A. N., DESBRUN M., SCHRÖDER P.: Discrete shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2003), SCA '03, Eurographics Association, p. 62–67. 6

[Gib10] GIBSON R. F.: A review of recent research on mechanics of multifunctional composite materials and structures. *Composite Structures 92*, 12 (2010), 2793–2810. 1

[GKS02] GRINSPUN E., KRYSL P., SCHRÖDER P.: CHARMS: a simple framework for adaptive simulation. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2002), SIGGRAPH '02, ACM, p. 281–290. 2

[GSS*15] GAST T. F., SCHROEDER C., STOMAKHIN A., JIANG C.,

TERAN J. M.: Optimization integrator for large time steps. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 21*, 10 (2015), 1103–1115. 3

[HZG*18] HU Y., ZHOU Q., GAO X., JACOBSON A., ZORIN D., PANOZZO D.: Tetrahedral meshing in the wild. *ACM Trans. Graph. 37*, 4 (2018). 2

[JSS*15] JIANG C., SCHROEDER C., SELLE A., TERAN J., STOMAKHIN A.: The affine particle-in-cell method. *ACM Trans. Graph. 34*, 4 (2015). 2

[JZH*21] JIANG Z., ZHANG Z., HU Y., SCHNEIDER T., ZORIN D., PANOZZO D.: Bijective and coarse high-order tetrahedral meshes. *ACM Trans. Graph. 40*, 4 (July 2021). 2

[LKJ21] LI M., KAUFMAN D. M., JIANG C.: Codimensional incremental potential contact. *ACM Trans. Graph. 40*, 4 (2021). 1

[LLK*20] LONGVA A., LÖSCHNER F., KUGELSTADT T., FERNÁNDEZ-FERNÁNDEZ J. A., BENDER J.: Higher-order finite elements for embedded simulation. *ACM Trans. Graph. 39*, 6 (Nov. 2020). 2

[Llo82] LLOYD S.: Least squares quantization in PCM. *IEEE Transactions on Information Theory 28*, 2 (1982), 129–137. 5

[MGL*15] MALGAT R., GILLES B., LEVIN D. I. W., NESME M., FAURE F.: Multifarious hierarchies of mechanical models for artist assigned levels-of-detail. In *Proceedings of the 14th ACM SIGGRAPH / Eurographics Symposium on Computer Animation* (New York, NY, USA, 2015), ACM, pp. 27–36. 1, 2, 4

[MKB*10] MARTIN S., KAUFMANN P., BOTSCH M., GRINSPUN E., GROSS M.: Unified simulation of elastic rods, shells, and solids. *ACM Trans. Graph. 29*, 4 (2010). 1, 2

[MMCK14] MACKLIN M., MÜLLER M., CHENTANEZ N., KIM T.-Y.: Unified particle physics for real-time applications. *ACM Trans. Graph. 33*, 4 (2014). 2

[MTGG11] MARTIN S., THOMASZEWSKI B., GRINSPUN E., GROSS M.: Example-based elastic materials. *ACM Trans. Graph. 30*, 4 (2011). 3

[MWN*17] MANTEAUX P.-L., WOJTAN C., NARAIN R., REDON S., FAURE F., CANI M.-P.: Adaptive physically based models in computer graphics. *Computer Graphics Forum 36*, 6 (2017), 312–337. 1, 2

[NSO12] NARAIN R., SAMII A., O'BRIEN J. F.: Adaptive anisotropic remeshing for cloth simulation. *ACM Trans. Graph. 31*, 6 (2012), 152:1–152:10. 1, 2

[OH99] O'BRIEN J. F., HODGINS J. K.: Graphical modeling and animation of brittle fracture. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques* (USA, 1999), SIGGRAPH '99, ACM Press/Addison-Wesley Publishing Co., p. 137–146. 6

[RCPO22] ROMERO C., CASAS D., PÉREZ J., OTADUY M.: Contact-centric deformation learning. *ACM Trans. Graph. 41*, 4 (2022). 2

[SBO18] SÁNCHEZ-BANDERAS R. M., OTADUY M. A.: Strain rate dissipation for elastic deformations. *Computer Graphics Forum 37*, 8 (2018), 161–170. 4

[SDG*19] SCHNEIDER T., DUMAS J., GAO X., BOTSCH M., PANOZZO D., ZORIN D.: Poly-spline finite-element method. *ACM Trans. Graph. 38*, 3 (Mar. 2019). 2

[SGK18] SMITH B., GOES F. D., KIM T.: Stable neo-hookean flesh simulation. *ACM Trans. Graph. 37*, 2 (2018). 7

[SHD*18] SCHNEIDER T., HU Y., DUMAS J., GAO X., PANOZZO D., ZORIN D.: Decoupling simulation accuracy from mesh quality. *ACM Trans. Graph. 37*, 6 (Dec. 2018). 2

[SYS*21] SHEN S., YANG Y., SHAO T., WANG H., JIANG C., LAN L., ZHOU K.: High-order differentiable autoencoder for nonlinear model reduction. *ACM Trans. Graph. 40*, 4 (2021). 2

[WDGT01] WU X., DOWNES M. S., GOKTEKIN T., TENDICK F.: Adaptive Nonlinear Finite Elements for Deformable Body Simulation Using Dynamic Progressive Meshes. *Computer Graphics Forum* (2001). 2

[ZBLJ20] ZHANG J. E., BANG S., LEVIN D. I. W., JACOBSON A.: Complementary dynamics. *ACM Trans. Graph. 39*, 6 (2020). 2