



SMI 2012: Short Paper

Mixed-element volume completion from NURBS surfaces[☆]Tobias Martin^{a,c,*}, Elaine Cohen^a, Robert M. Kirby^b^a School of Computing, University of Utah, United States^b Scientific Computing and Imaging Institute, University of Utah, United States^c ETH Zürich, Switzerland

ARTICLE INFO

Article history:

Received 3 December 2011

Received in revised form

15 March 2012

Accepted 17 March 2012

Available online 28 March 2012

Keywords:

Model completion

Volumetric parameterization

Meshing

Physically-based animation and simulation

ABSTRACT

In this paper we present a methodology to create a volumetric representation from a 2-manifold without boundaries represented with untrimmed NURBS surfaces. A trivariate NURBS representation is difficult to construct from such a representation, especially when the input surface patches were created with only a boundary representation as the goal. These kinds of inputs arise in numerous existing geometric modeling scenarios such as models from CAD systems, subdivision surfaces, quadrilateral meshing and data-fitting. Our approach, nearly automatic and only requiring minimal user input, creates a hybrid representation using trivariate NURBS elements at the boundary and unstructured tetrahedral elements in the interior of the object. The original boundary representation of the input model is maintained in the final representation, allowing the volumetric representation to be used in both computer graphics simulations and isogeometric analysis applications. We demonstrate that the hybrid representation yields convergence under model refinement, and that it can be used efficiently in elastic body simulation.

© 2012 Elsevier Ltd. All rights reserved.

1. Introduction

Objects represented with NURBS surfaces are widely used in computer graphics and engineering applications. A subdivision surface is frequently converted into a set of NURBS surfaces for rendering purposes, and mechanical parts are modeled or reverse-engineered using NURBS patches of higher-order polynomial degrees. In the isogeometric analysis paradigm [1], smooth basis functions used to model the smooth geometry are also used as basis for the simulation. The analysis result is stored as an attribute of the model representation. This avoids both the need to generate a finite element mesh and the need to reverse-engineer a CAD model from the finite element mesh. However, a volumetric representation must be generated if a simulation methodology is to be applied to the volume of the input.

If the input consists of NURBS surface patches and if its parameterization must be maintained in the resulting volumetric representation, it is a difficult problem to complete its volume with trivariate NURBS elements, and has not been solved for arbitrary shapes models. In practice, a volumetric representation is constructed that either consists of (1) classic finite elements (e.g., [2,3]) or (2) larger smooth trivariate B-spline patches. Methods in category (1) are mostly automatic. Methods in

category (2) (e.g., [4–10]) are generally based on a volumetric parameterization established on the volume from which the volumetric representation is constructed and may require user interaction. While in both categories, a full volumetric hexahedral representation is created, the original higher-order representation is not part of the simulation representation. Therefore, the isogeometric analysis paradigm is not fully implemented.

This paper makes three contributions: (1) Presentation of a new mixed volumetric representation that keeps the original NURBS surface representation and parameterization. (2) Creation of a semi-structured NURBS volume with arbitrary thickness as a user parameter (with a few restrictions), where the remaining interior volume is completed with unstructured tetrahedra (Fig. 1). (3) A collocation approach which unifies the different element types and verification of the representation in a simulation setting.

The proposed approach, based on harmonic functions [11] and a sampled midstructure, allows offsetting the input model into its interior without creating degenerate elements to produce a volumetric semi-structured NURBS volumetric mesh offset inward from the boundary of the object. It significantly reduces the time needed to create a volumetric model and avoids the correspondence problem by matching the interior boundary NURBS elements with linear tetrahedra.

NURBS elements require more computational effort during simulation, e.g., numerical integration and evaluation of simulation quantities due to higher-order polynomial degree. These quantities can be efficiently evaluated on linear tetrahedra due to its lower degree. However, simulation applied to structured NURBS elements can produce simulation results of higher quality

[☆] If applicable, supplementary material from the author(s) will be available online after the conference. Please see <http://dx.doi.org/10.1016/j.cag.2012.03.008>.

* Corresponding author.

E-mail address: martin@cs.utah.edu (T. Martin).

because of smoothness across elements [1]. Fig. 2 illustrates that the proposed representation can be useful in applications such as animating elasticity, where high quality NURBS elements on the domain’s boundary prevent element degeneracies. Linear elements, prone to flipping in non-linear simulation scenarios, are placed away from the critical regions, and the overall impression of the animation visually appears smoother as well.

Related work is discussed in Section 2, the hybrid representation is introduced in Section 3. Section 4 verifies convergence of the hybrid representation under refinement. The paper discusses results in Section 5 and concludes in Section 6.

2. Related work

Hybrid meshing or mixed element methods, where prisms or hexahedra are combined with tetrahedra, have been used in various engineering scenarios, e.g., aerospace applications [12], geophysics applications and computation fluid dynamics applications [13]. If the input surface is a quadrilateral mesh, pyramids are used to transition from boundary hexahedra to tetrahedra. H-morph [14], based on an advancing front scheme, iteratively converts an input tetrahedral mesh into a hexahedral mesh. Tetrahedral elements that cannot be converted into hexahedral elements are connected to their adjacent hexahedral elements using pyramids.

In these cases, C^0 continuity between two element types is maintained. Our approach does not enforce such continuity and in that respect is similar to, and motivated by approaches discussed in the introduction of Section 3.4, where the nodes defining the linear tetrahedra are linked to these higher-order elements. We show that the resulting representation yields convergence under

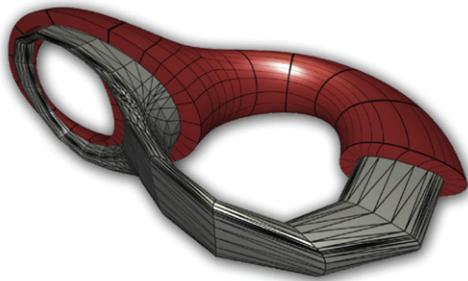


Fig. 1. Hybrid representation: tri-cubic NURBS elements (red) at the boundary. Linear tetrahedra (gray) in the interior of the object. The surface patches are constructed using Catmull–Clark subdivision. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

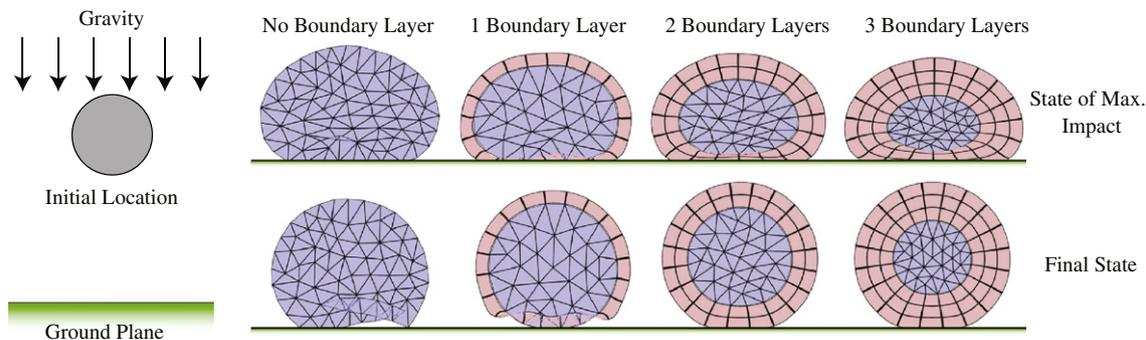


Fig. 2. Effect of C^2 elements at domain boundary: the configuration of this 2D elasticity example, based on a Hookean material using Chauchy–Green strain, is illustrated on the left. The top row shows the state of maximal impact, i.e., the state where the elastic potential energy reaches its maximum, when the shape hits the ground plane. The bottom row shows the final state of the object. By increasing the number of boundary layers and its thicknesses, the corresponding final states do not exhibit flipped elements. Note, due to the gravitational force the final shape is not circular.

mesh refinement. It will also be shown that our hybrid approach has a similar convergence rate to a representation which uses only tetrahedral elements, but has the added benefit that it can avoid some of the stability problems of tetrahedra only meshes.

In a similar vein to our work is the method in [15] generating a hybrid mesh consisting of prisms and tetrahedra. The hybrid representation is constructed from a triangulated surface representing the domain of interest. This triangulated surface is offset along a marching direction. A marching vector of a node is determined through local geometric constraints and a marching step size is chosen to reduce curvature from that of the previous marching surface. Similarly, [16] uses solutions to the Eikonal equation. Finally, tetrahedra are used to fill up the rest of the object. More recently, [17] first computes a distance field from an input triangle mesh. Based on a user specified thickness parameter, an isosurface of the distance field is extracted and a hexahedral shell mesh based on a polycube domain [18] and harmonic functions is constructed, i.e., the remaining volume in the object is void. Peng et al. [19] proposes a method to add volume textures to an input surface. The method offsets the input surface based on an ODE. The ODE is designed such that the limit offset surface is the medial axis of the input surface. Note, when the input surface contains sharp edges the medial axis touches the input surface, and the outer layer is very thin at the sharp edges.

In our approach, harmonic functions are used in combination with a midstructure representation to define the offset function for the exterior surface. The flexibility of harmonic functions allow the user to specify any appropriate midstructure, e.g., a point-sampled simplified medial axis or a 1D curve skeleton of the object or an isosurface as used in [17]. Also, marching directions and step sizes are implicitly determined and guarantee not to introduce degenerate elements because of the maximum principle of harmonic functions. Furthermore, compared to normal offsetting, harmonic functions allow offsetting the input surface deeper into the interior of the object. The remaining volume is completed with tetrahedra.

3. Volumetric representation

This section describes our proposed modeling pipeline. Let the domain of interest be the input \mathcal{S} , a 2-manifold without boundaries. \mathcal{S} consists of a collection of coefficients $c_l \in \mathbb{R}^3$ and associated weights $h_l \in \mathbb{R}$ and a set of parametric patches s_k , defined as

$$s_k(u, v) := \frac{\sum_{i=0}^p \sum_{j=1}^q h_{ij} c_{ij} B_{i,p,\tau_u}(u) B_{j,q,\tau_v}(v)}{\sum_{i=0}^p \sum_{j=1}^q h_{ij} B_{i,p,\tau_u}(u) B_{j,q,\tau_v}(v)}, \quad (1)$$

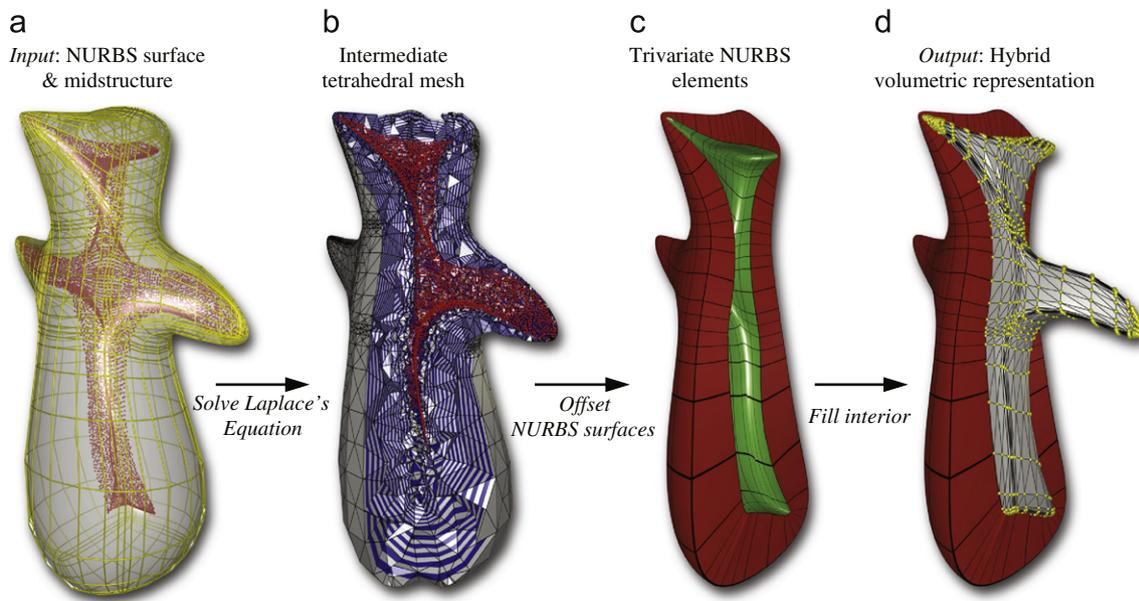


Fig. 3. The input are NURBS surface patches, in these cases generated from a Catmull–Clark subdivision model. A sampled midstructure such as a simplified medial axis is used to generate trivariate NURBS patches at the boundary of the volumetric representation. The remaining interior is filled up with unstructured tetrahedral elements. (Elements in (b)–(d) are culled to visualize the interior.)

where $B_{i,p,\tau_u}(u)$ and $B_{j,q,\tau_v}(v)$ are B-spline basis functions as defined in [20], p and q the degrees and τ_u and τ_v the local knot vectors of patch $s_k(u,v)$. The coefficients c_{ij} with corresponding scalar weights h_{ij} define a control mesh of dimension $(p+1) \times (q+1)$. Note that the indices (i,j) in Eq. (1) are local to the patch $s_k(u,v)$, i.e., there exists a mapping to a global index l .

S and an associated sampled midstructure (generated by using one of several methods, e.g., [21]) are used to construct a volumetric representation in four steps, as shown in Fig. 3.

- Step 1. Triangulate control mesh and create unstructured tetrahedral mesh containing the point-sampled midstructure, Section 3.1
- Step 2. Construct a harmonic function on the tetrahedral mesh (Fig. 3b), Section 3.1
- Step 3. Offset the input surface using the harmonic function, allowing the user to specify the thickness of the resulting semi-structured volumetric NURBS representation (Fig. 3c), Section 3.2
- Step 4. Triangulate the limit surface (green surface in Fig. 3c) which has the same mesh layout as the input, and fill its interior with unstructured tetrahedral elements (Fig. 3d), Section 3.4

The general definition of S and its patches $s_k(u,v)$ allow a wide range of input representations. For instance, if the input is an unstructured quadrilateral mesh, then $s_k(u,v)$ defines a bi-linear patch with $p=q=1$ and $\tau_u = \tau_v = \{0,0,1,1\}$ and $h_l=1$. A more general input may use different weights and a mix of floating and open knot vectors (see [20]), allowing higher continuity among adjacent patches and the representation of rational geometries.

3.1. Harmonic mapping

The first step in our framework consists of triangulating the set of coefficients c_l by triangulating the regular local control mesh c_{ij} for each patch $s_k(u,v)$ as defined in Eq. (1). From this triangle mesh and a sampled midstructure of S , an unstructured tetrahedral mesh \mathcal{H} is constructed (e.g., by using [3]). Let \mathcal{V} be the set of vertices in \mathcal{H} . Let \mathcal{V}_E be the set of vertices on the boundary of \mathcal{H} and \mathcal{V}_I be the set of

vertices defining the midstructure of S lying in the interior of \mathcal{H} . Note, \mathcal{V}_E and \mathcal{V}_I are subsets of \mathcal{V} and also that \mathcal{V}_E may contain vertices in addition to the coefficients c_l , depending on the tetrahedralization.

Given \mathcal{H} , \mathcal{V}_E and \mathcal{V}_I , a scalar function w on \mathcal{H} is computed such that it satisfies Laplace's equation,

$$\nabla^2 w = 0. \quad (2)$$

w is harmonic and satisfies the maximum principle, and therefore has been used in the domain of meshing and volumetric parameterization (e.g., [7,17]).

The finite element method [22] is used to discretize Eq. (2). Vertices in \mathcal{V}_E are set to 0, vertices in \mathcal{V}_I are set to 1 (Dirichlet BC). With this setup $w(x,y,z)$ evaluates to 0 at the boundary of \mathcal{H} and to 1 at vertices defining the midstructure (see Fig. 3b). A solution has the form $w(x,y,z) = \sum_{\mathbf{v}_k \in \mathcal{V}} \hat{w}_k \phi_k(x,y,z)$, where $\phi_k(x,y,z)$ are linear hat functions [22] associated with the respective vertices in \mathcal{H} . Galerkin's formulation [22] is used to set up a linear system which is solved to assign a harmonic value with every vertex in \mathcal{V} . The gradient field ∇w over \mathcal{H} is piece-wise constant, flowing towards the sampled midstructure.

3.2. Offset input surface

The harmonic function $w(x,y,z)$ is used to offset the vertices c_l defining S into the interior of S . In a first step, the user is given the ability to choose a parameter value $w_0 \in [0,1]$, allowing the user to control the volume enclosed by S and the isosurface $w(x,y,z) = w_0$ (see Fig. 3c). Once this choice has been made, the remaining pipeline stages proceed automatically. Note that since the midstructure representation is sampled, the isosurface at w_0 could be of higher genus than the input S , especially when w_0 is close to 1. The proposed method requires that the isosurface at w_0 and the input has the same genus, i.e., the user must choose a w_0 that satisfies this condition.

For every c_l defining the input surface S a path $g_l(\omega)$, $g_l: \mathbb{R} \rightarrow \mathbb{R}^3$, is constructed emanating from c_l by following ∇w and terminating at point \bar{c}_l , where $w(\bar{c}_l) = w_0$, i.e., $g_l(0) = c_l$ and $g_l(w_0) = \bar{c}_l$. Note that $g_l(\omega)$ is parameterized through $w(x,y,z)$. Let \bar{S} be the surface defined with the coefficients \bar{c}_l (see green surface in Fig. 3c), with corresponding surface patches $\bar{s}_k(u,v)$.

Given the paths $g_l(\omega)$, surface patches $s_k(u, v)$ can be swept into the interior of \mathcal{S} . Assume there exists a sorted set $W = \{\omega_0 = 0, \dots, \omega_n = w_0\}$, consisting of n scalars, where $\omega_i \in [0, w_0]$. Given W , for each $s_k(u, v)$ a volumetric NURBS patch $v_k(u, v, w)$, defined as

$$v_l(u, v, w) := \sum_{i,j,k=0}^{p,q,n} w_{i,j,k} g_{i,j}(\omega_k) \mathcal{B}_{i,j,k}(u, v, w) \quad (3)$$

can be constructed, where

$$\mathcal{B}_{i,j,k} := B_{i,p,\tau_u}(u) B_{j,q,\tau_v}(v) B_{k,r,\tau}(w), \quad (4)$$

and where $v_l(u, v, w)$ defines a triple sum. r defines the degree in w and τ is the knot vector in the domain $[0, 1]$. The weights $w_{i,j,k}$ in the interior of \mathcal{V} are set to 1. W is determined through an optimization procedure trying to determine the values ω_i such that the volumetric elements in $v_l(u, v, w)$ have equal volume.

Due to the maximum principle of w , paths $g_l(\omega)$ are consistent and therefore self-intersecting elements, as they often arise by offsetting a surface along its normal field, generally do not occur.

Let \mathcal{V} be the collection of volumetric patches $v_l(u, v, w)$. Note that the outer boundary of \mathcal{V} and the input surface \mathcal{S} are the same. The inner boundary of \mathcal{V} , i.e., $\bar{\mathcal{S}}$, (see red surface in Fig. 4) is geometrically similar to the isosurface $w(x, y, z) = w_0$.

3.3. Tetrahedralize the interior

For the following discussion let \mathcal{I} be a collection of tuples $(\bar{c}_l, t_l, \bar{s}_k(u, v))$, where $t_l = (u_l^*, v_l^*)$ is the node location [20] corresponding to \bar{c}_l , and $\bar{s}_k(u, v)$ is the surface patch whose parametric domain contains t_l . Given such a tuple, let $x_l = \bar{s}_k(t_l)$. x_l is equivalent to a first-order projection of \bar{c}_l onto $s_k(u, v)$ [20] and can be computed efficiently. Fig. 5 illustrates this setup for a single NURBS surface and corresponding triangle mesh.

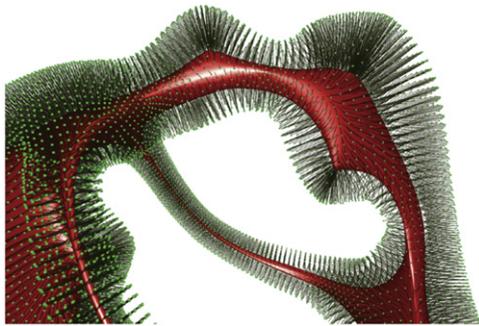


Fig. 4. Paths following ∇w (gray) emanate at c_l (green points) and terminate on isosurface $w(x, y, z) = w_0$, approximated with $\bar{\mathcal{S}}$ (red). The choice of the midstructure allows creation of elements in thinner tubular regions. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

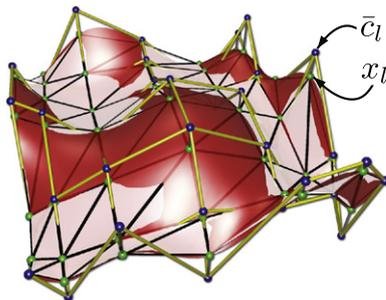


Fig. 5. Illustration of C^2 NURBS surface and corresponding C^0 triangle mesh. \bar{c}_l (blue) for the NURBS surface and x_l (green) for the triangle mesh boundary represent the same degree of freedom. Due to geometric concavities, the representations are overlapping.

The final step in our proposed framework consists of filling up the interior of $\bar{\mathcal{S}}$ with unstructured tetrahedra. Similarly as in the first step of our framework (Section 3.1), $\bar{\mathcal{S}}$ is triangulated. However, instead of the coefficients \bar{c}_l , the locations x_l are used for the triangulation (see Fig. 5). The interior of the resulting triangle mesh is filled up with unstructured tetrahedra to construct the tetrahedral mesh \mathcal{T} (Fig. 3d).

3.4. Hybrid representation

After executing these framework steps, the proposed hybrid volumetric representation is defined by the tuple

$$\mathcal{H} = (\mathcal{V}, \mathcal{T}, \mathcal{I}). \quad (5)$$

Even though \bar{c}_l and its corresponding x_l in \mathcal{I} are at different locations (Figs. 2 and 6), in our simulation framework discussed below, they represent the same degree of freedom. \mathcal{I} is the interface between \mathcal{V} and \mathcal{T} .

The connection of two different representations of interfaces, in our case \mathcal{V} and \mathcal{T} , has been examined already in the context of mechanics modeling, and in the areas of contact and fluid structure interaction (see e.g., [23–26]). Our proposed approach has a similar setup as a mortar approach in the sense that separate finite element discretizations on nonoverlapping subdomains are used (Fig. 6).

However, our approach is more similar to a collocation approach, where \mathcal{V} and \mathcal{T} are forced to match at specified collocation points, i.e. the points in \mathcal{I} . Given this setup, it is clear that the two interfaces will not match because of the difference in representations (see Figs. 5 and 6). The interface will have gaps or overlaps depending on the concave or convex regions in $\bar{\mathcal{S}}$. In our case, we are interested in volumetric representations, so although overlapping volumetric representations are not the desired goal, we can be confident that the two interfaces (and hence the volumetric representations) converge to a common representation. This is due to B-spline properties [20], where under successive refinement, the control mesh of $\bar{\mathcal{S}}$ converges to the surface $\bar{\mathcal{S}}$ and hence the gaps and overlaps between the boundary of \mathcal{T} and $\bar{\mathcal{S}}$ get smaller. This is in contrast to NEFEM, proposed in [27]. NEFEM uses straight sided triangles internally and triangles with one isoparametric curved edge for those that have an edge on the boundary so it maintains the exact NURBS input boundary in the output representation. For those boundary elements, a specifically designed numerical integration is presented. The solution space is based on polynomials that are C^0 across element edges.

This mixed representation choice has been made for several reasons. While continuity could be more easily achieved in the 2D case, the proposed method avoids patching higher-order tetrahedral elements to quadrilateral surface patches in 3D. This choice also simplifies model subdivision and generation of \mathcal{T} . Secondly, in Section 4 it will be demonstrated that simulation

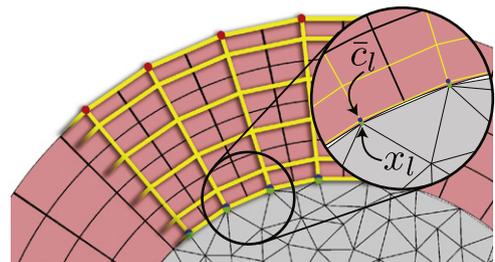


Fig. 6. Hybrid representation in 2D. Bi-quadratic NURBS elements were chosen at the boundary to demonstrate the worst-case scenario of the alignment with triangles in the interior. Convergence in simulations achieved even in this more general scenario.

convergence on a 2D problem is achieved. Furthermore, the accompanied video shows that dynamic physically-based animation is efficiently applied to the 3D case.

Given a static or dynamic problem, discretized using the finite element method [22], in the general case, there exists a solution function $\alpha : \Omega \rightarrow \mathbb{R}^d$. For instance in linear elasticity $d=3$ where α describes a displacement field defined over Ω , or in heat conduction where α represents a temperature scalar field defined over Ω , i.e., $d=1$. In our framework Ω is represented with \mathcal{H} .

Given a tuple $(\bar{c}_l, t_l, \bar{s}_k(u, v)) \in \mathcal{I}$. Let α_l be the coefficient corresponding to the vertex x_l of a tetrahedron in \mathcal{T} . Note that α evaluates to α_l at this tetrahedron at x_l , i.e., $\alpha_l := \alpha(x_l)$. Let α_s be the evaluation of α at $\bar{s}_k(t_l)$, i.e., $\alpha_s := \alpha(\bar{s}_k(t_l))$. Given the discontinuous nature of \mathcal{H} , $\alpha_l \neq \alpha_s$ even though $x_l = \bar{s}_k(t_l)$. Therefore, before simulation proceeds, α_s is assigned to the coefficient α_l . This enforces that $\alpha(\bar{s}_k(t_l))$ and $\alpha(x_l)$ match. In Section 4 we demonstrate that convergence can be achieved with the proposed collocation-based approximation.

3.5. Robustness and practical considerations

\bar{s} tends to have self-intersections if the input surface contains highly stretched elements, e.g., see Fig. 1. Therefore, before the hybrid representation generation framework is executed, the input surface is appropriately refined, such that elements have similar size and shape. While the element count of the input surface is increased, the resulting triangle mesh corresponding to \bar{s} has a better quality and is more suitable to generate a tetrahedral mesh. Furthermore, while paths computed from a harmonic field are guaranteed to be free of intersections, paths traced on a discretized field can overlap. The resulting higher-order elements comprising the thick shell can therefore have self-intersections which should be avoided. Similarly to the method in [7], paths are traced simultaneously in small steps, where in each step the front defined by the path endpoints are smoothed using a Laplacian smoothing scheme such that endpoints remain on the corresponding isosurface.

4. Convergence study in 2D

Here, we examine a study in 2D to show that our proposed hybrid representation as discussed in Section 3.4 produces comparable simulation results with respect to convergence rates under h -refinement [22] as the equivalent representation which uses only triangles.

We are given the smooth function $g : \Omega \rightarrow \mathbb{R}$

$$g(x, y) := J(4, J_0(4, 2)r(x, y))\sin(4\theta(x, y)), \quad (6)$$

where $r(x, y) := \sqrt{x^2 + y^2}$, and $\theta(x, y)$ defines the angle between vector (x, y) and the Cartesian coordinate axes, i.e., $r(x, y)$ and $\theta(x, y)$

convert (x, y) into polar coordinates. Ω in this study represents a disk centered at the origin with a radius of 1 with boundary $\partial\Omega$. Furthermore, $J(n, z)$ is the n th Bessel function of the first kind at $z \in \mathbb{R}$, and $J_0(n, m)$ is the m th zero of the n th Bessel function of the first kind. Since $g(x, y) = 0$ at $\partial\Omega$, the Dirichlet boundary condition at $\partial\Omega$ is set to zero. In the following experiment, let $f(x, y) := \nabla^2 g$.

In this study we investigate Poisson's equation $-\nabla^2 \tilde{g} = f$, solved using Galerkin's method on two disk representations: (1) the disk is represented with NURBS elements at the boundary and triangles in the interior using the hybrid framework as discussed in Section 3.4; (2) the disk is represented with triangles only. Both representations are refined three times where the error $\|g(x, y) - \tilde{g}(x, y)\|_2$ is computed for each refinement step. The corresponding log-log diagram is shown in Fig. 7. The figure also shows the disk representations at an intermediate refinement step and the corresponding solution along the z -axis of the respective disk.

The convergence of a triangle representation to the true solution is h^{p+1} , where p is the order on element and h its radius. For instance, quadratic NURBS elements therefore have a cubic convergence rate [28], while linear tetrahedral elements converge quadratically. However, as the experiment in this section shows, under h -refinement, the slope for both representations is approximately -2 indicating quadratic convergence even in the presence of higher-order NURBS elements. Our hybrid representation has two inherent approximation errors: (1) geometric approximation error of the interior representation and (2) approximation error of the field. Since both geometric subdivision and linear triangle elements have a quadratic convergence rate, the overall convergence rate is quadratic as well. We did not study the effects from varying depths of the NURBS representation and how convergence is affected. This is subject to future research.

5. Results

We applied the proposed approach to the following closed input objects, represented with the bi-cubic patches: B-spline surfaces computed from Catmull-Clark subdivision surfaces shown in Figs. 1 and 3; Bézier surfaces with G^1 continuity computed from quadrilateral meshes computed using the approach by [29] shown in Fig. 8(a) and (b); NURBS surfaces approximating a mechanical part with C^2 continuity in the interior and C^0 at adjacent surface patches (Fig. 8(c)). In all cases, a cubic knot vector is used for the w -direction. In the current framework, the midstructure has been computed using the tight co-cone approach [21], where for some models, the resulting medial axis has been manually simplified. For instance, for the fandisk model (Fig. 8(c)), sheets corresponding to sharp features were manually removed. Although manual medial axis simplification is a time-consuming procedure, our approach only requires a sampling of a midstructure and does not require

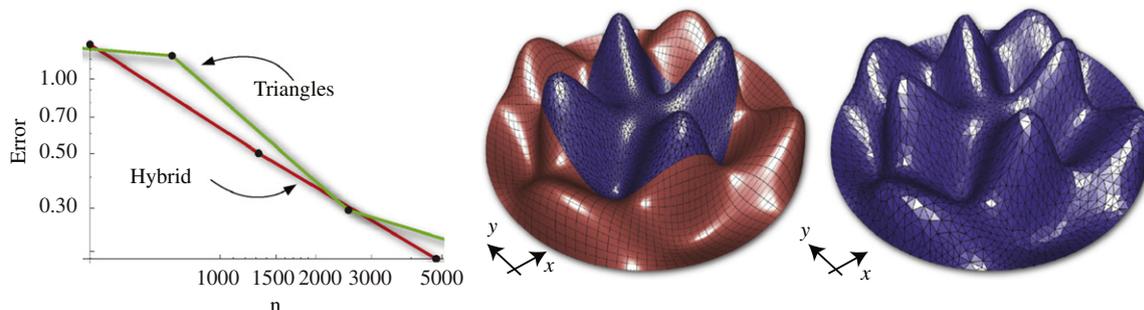


Fig. 7. Left: convergence plot comparing the hybrid representation with the representation which only uses triangles. For the error computation, we chose the l_2 norm between the approximated and true solution. n is the degree of freedoms for the refinement steps. Solutions along z -axis on hybrid representation (middle) and triangle only representation (right), during a h -refinement step.

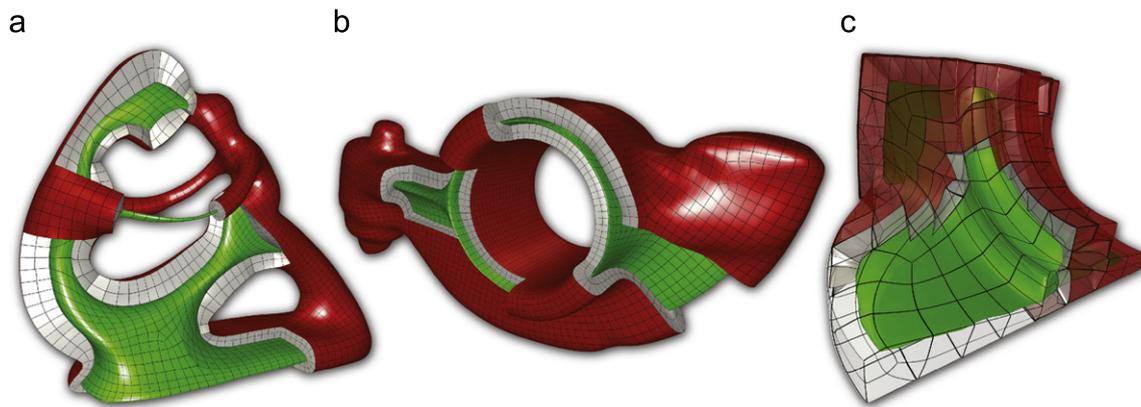


Fig. 8. Methodology applied to three objects. Object boundaries are shown in red, interior boundaries in green. The input for (a) and (b) are bi-cubic Bézier patches, and for (c) bi-cubic NURBS surfaces. The wire-frame shows C^0 continuities. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

proper topological connectivity, so the simplification was performed rapidly and did not exceed 10 min for each of the input models. We applied Laplacian smoothing and edge collapses to produce triangles of similar size and area. The corresponding vertex set was used as sampling for the intermediate tetrahedral mesh (Fig. 3(b)).

The only remaining user input to the subsequent pipeline stages was the choice of thickness of the outer B-spline volumetric region as discussed in Section 3.1. The reason for this is that a user might want to specify a different thickness depending on the application. For instance, a thicker NURBS volume would result in more high-order trivariate patches with wider basis function support, resulting in slightly slower simulation, but higher simulation quality as illustrated in Fig. 2. However, for prototyping purposes, a user might choose thinner volumes to compute simulation results more quickly.

The remaining computational step consisting of constructing Laplace's matrix, tracing paths and creating volumetric elements, can be neglected as these operations can be parallelized and computed efficiently. Timings were less than 5 min for each of the test models. As discussed in Section 3.2, after the user performed all required input, the sampling procedure in [7] is adapted for our context that resamples the paths to create roughly equal B-spline volumes by avoiding degenerate elements. Note that this procedure only re-samples the paths $g_i(\omega)$ computed from the input points and ∇w as discussed in Section 3.2 and does not require any other re-computation.

6. Conclusion

In this paper we proposed a framework to create a volumetric representation from an input surface represented with untrimmed NURBS surfaces from various inputs, making the creation of a volumetric representation independent from creating the input surfaces. The approach is based on harmonic functions which are used to offset the input surface until a user-specified offset distance is reached. The remaining interior volume is filled with unstructured tetrahedra. A collocation approach is proposed to address the correspondence problem. The boundary surface of the volumetric representation and the input surface are the same, i.e., smoothness and geometry of the input surface is maintained in the volumetric representation.

A limitation of the approach is the geometric discontinuity between the high-order boundary in the interior and the boundary of the interior tetrahedral mesh, where the convergence rate of the tetrahedral mesh is the limiting factor. We have demonstrated that the resulting hybrid volumetric representation is

stable, can be used in simulation scenarios, and is motivated by discontinuous hybrid simulation representations from other areas. In future work, we plan to extend our framework so that it can be applied to a wider range of element types, for instance higher-order tetrahedral elements. Note that a fully continuous representation is more difficult to achieve, mainly because of the smooth input surface representations which can have an arbitrary choice of continuity and smoothness properties.

In conclusion, we see our approach especially useful in applications which require that the input surface parameterization must be maintained in the volumetric representation, such as in isogeometric analysis. But the approach is also useful in applications where the time to generate volumetric inputs is limited, requiring only little user input, which is generally desired and even required in various computer graphics applications.

Acknowledgments

This publication is supported in part by NSF IIS-1117997 and ARO W911NF-08-0517 (Program Manager Dr. Mike Coyle). Various models in the paper are provided by the AIM@SHAPE Shape Repository.

Appendix A. Supplementary material

Supplementary data associated with this article can be found in the online version of <http://dx.doi.org/10.1016/j.cag.2012.03.008>.

References

- [1] Hughes TJR, Cottrell JA, Bazilevs Y. Isogeometric analysis: Cad, finite elements, nurbs, exact geometry, and mesh refinement. *Comput Methods Appl Mech Eng* 2005;194:4135–95.
- [2] Owen SJ. A survey of unstructured mesh generation technology. In: *International Meshing Roundtable*; 1998. p. 239–67.
- [3] Si H. Tetgen: a quality tetrahedral mesh generator and three-dimensional Delaunay triangulator; 2005 <<http://tetgen.berlios.de>>.
- [4] Wang H, Jin M, He Y, Gu X, Qin H. User-controllable polycube map for manifold spline construction. In: *SPM '08: proceedings of the 2008 ACM symposium on solid and physical modeling*. New York, NY, USA: ACM; 2008. p. 397–404.
- [5] He Y, Wang H, Fu C-W, Qin H. Technical section: a divide-and-conquer approach for automatic polycube map construction. *Comput Graph* 2009;33(3):369–80.
- [6] Li B, Li X, Wang K, Qin H. Generalized polycube trivariate splines. In: *Shape modeling international conference (SMI)*; 2010. p. 261–5. <http://dx.doi.org/10.1109/SMI.2010.40>.

- [7] Martin T, Cohen E, Kirby RM. Volumetric parameterization and trivariate b-spline fitting using harmonic functions. *Comput Aided Geom Des* 2009; 26:648–64.
- [8] Martin T, Cohen E. Volumetric parameterization of complex objects by respecting multiple materials. *Comput Graph* 2010;34(3):187–97.
- [9] Nieser M, Reitebuch U, Polthier K. CUBECOVER—parameterization of 3d volumes. *Comput Graph Forum* 2011;30(5):1397–406.
- [10] Gregson J, Sheffer A, Zhang E. All-hex mesh generation via volumetric polycube deformation. *Comput Graph Forum* 2011;30(5):1407–16.
- [11] Ni X, Garland M, Hart JC. Fair morse functions for extracting the topological structure of a surface mesh. In: *Proceedings of the SIGGRAPH; 2004*.
- [12] Khawaja A, Kallinderi Y. Hybrid grid generation for turbomachinery and aerospace applications. *Int J Numer Methods Eng* 2000;49(1):145–66.
- [13] Ito Y, Shih AM, Soni BK. Hybrid mesh generation with embedded surfaces using a multiple marching direction approach. *Int J Numer Methods Fluids* 2011;67(1):1–7. <http://dx.doi.org/10.1002/flid.1962>.
- [14] Owen SJ, Saigal S. H-morph: an indirect approach to advancing front hex meshing. *Int J Numer Methods Eng* 2000;49(1–2):289–312.
- [15] Kallinderis Y, Khawaja A, McMorris H. Hybrid prismatic/tetrahedral grid generation for complex geometries. *AIAA Paper* 1996;34:93–0669.
- [16] Wang Y, Murgie S. Hybrid mesh generation for viscous flow simulation. In: Pbay PP, editor. *Proceedings of the 15th international mesh roundtable*. Springer-Verlag; 2006. p. 109–26.
- [17] Han S, Xia J, He Y. Hexahedral shell mesh construction via volumetric polycube map. In: *Proceedings of the 14th ACM symposium on solid and physical modeling, SPM '10*. New York, NY, USA: ACM; 2010. p. 127–36.
- [18] Tarini M, Hormann K, Cignoni P, Montani C. Polycube-maps. In: *SIGGRAPH '04: ACM SIGGRAPH 2004 papers*. New York, NY, USA: ACM; 2004. p. 853–60.
- [19] Peng J, Kristjansson D, Zorin D. Interactive modeling of topologically complex geometric detail. *ACM Trans Graph* 2004;23:635–43.
- [20] Cohen E, Riesenfeld RF, Elber G. *Geometric modeling with splines: an introduction*. Massachusetts, USA: A.K. Peters, Ltd.; 2001.
- [21] Dey TK, Goswami S. Tight cocone: a water-tight surface reconstructor. In: *SM '03: Proceedings of the eighth ACM symposium on solid modeling and applications*. New York, USA: ACM; 2003. p. 127–34.
- [22] Hughes TJR. *The finite element method: linear static and dynamic finite element analysis*. Dover; 2000.
- [23] Guillard H, Farhat C. On the significance of the geometric conservation law for flow computations on moving meshes. *Comput Methods Appl Mech Eng* 2000;190(11–12):1467–82.
- [24] Maday APY, Mavriplis C. Non-conforming mortar element methods: applications to spectral discretizations. *Domain Decomposition Methods, SIAM*.
- [25] Bernardi APC, Maday Y. A new nonconforming approach to domain decomposition: the mortar element method. In: Brezis H, Lions JL, editors. *Non-linear partial differential equations and their applications*. Pitman/Wiley.
- [26] Kirby RM, Yosibash Z, Karniadakis GE. Towards stable coupling methods for high-order discretization of fluid–structure interaction: algorithms and observations. *J Comput Phys* 2007;223:489–518.
- [27] Sevilla R, Fernández-Méndez S, Huerta A. NURBS-enhanced finite element method (NEFEM). *Arch Comput Methods Eng* 2011;18(4):441–84.
- [28] Bazilevs L, Beirão da Veiga Y, Cottrell J, Hughes T, Sangalli G. Isogeometric analysis: approximation, stability and error estimates for h-refined meshes. *Math Methods Models Appl Sci* 2006;16:1031–90.
- [29] Ray N, Li WC, Lévy B, Sheffer A, Alliez P. Periodic global parameterization. *ACM Trans Graph* 2006;25(4):1460–85.