

Supplemental: Higher-Order Finite Elements for Embedded Simulation

ANDREAS LONGVA, RWTH Aachen University

FABIAN LÖSCHNER, RWTH Aachen University

TASSILO KUGELSTADT, RWTH Aachen University

JOSÉ ANTONIO FERNÁNDEZ-FERNÁNDEZ, RWTH Aachen University

JAN BENDER, RWTH Aachen University

ACM Reference Format:

Andreas Longva, Fabian Löschner, Tassilo Kugelstadt, José Antonio Fernández-Fernández, and Jan Bender. 2020. Supplemental: Higher-Order Finite Elements for Embedded Simulation. *ACM Trans. Graph.* 39, 6, Article 181 (December 2020), 3 pages. <https://doi.org/10.1145/3414685.3417853>

1 CONVERGENCE RATE

Whereas measuring convergence rates for standard finite element discretizations is fairly straightforward, the measurement for embedded methods is more complicated by the need to design a measurement scheme which takes into account the precision with which the *embedded* domain is correctly captured. We discuss this problem more in detail in the following, and describe the setup we use to verify the convergence rate of the Finite Cell Method for a family of uniform background meshes with a non-trivial embedded geometry.

1.1 Measuring convergence rate

A popular and straightforward way to study convergence rates of finite element discretizations is to employ the *Method of Manufactured Solutions* (MMS) (see e.g. [Roache 2002]). The idea is to construct the problem from a prescribed exact solution – in this case a prescribed displacement function \mathbf{u} – by determining the remaining terms in the PDE so that they are consistent with the desired exact solution. By judiciously choosing a well-behaved, smooth, solution \mathbf{u} , optimal convergence rates in the L^2 and H^1 norms can then be attained. However, when manufacturing solutions in this way, one typically decides on the exact solution *before* deciding on the domain. Hence, the shape of the domain has no impact on the exact solution to the problem, and as a consequence, two different numerical solutions for two different domains will converge to the same solution when restricted to the intersection of the two domains. In the context of embedded methods, this means that solving the PDE on the background mesh with the standard FEM will essentially give the same solution as solving the PDE on the *exact* embedded domain. Hence, the MMS is not well-suited for studying the convergence rate of the FCM, as essentially the same solution is obtained regardless of whether the exact embedded domain is accurately captured or not.

Authors' addresses: Andreas Longva, longva@cs.rwth-aachen.de, RWTH Aachen University; Fabian Löschner, loeschner@cs.rwth-aachen.de, RWTH Aachen University; Tassilo Kugelstadt, kugelstadt@cs.rwth-aachen.de, RWTH Aachen University; José Antonio Fernández-Fernández, fernandez@cs.rwth-aachen.de, RWTH Aachen University; Jan Bender, bender@cs.rwth-aachen.de, RWTH Aachen University.

© 2020 Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3414685.3417853>.

More than just confirming expected convergence rates, we would like to contrast them with the results of simply applying the standard FEM applied to the background mesh. This leaves us with some constraints on how we can formulate our test problem:

- The domain should contain some flat region on which homogeneous Dirichlet boundary conditions (zero displacement) can be consistently applied across all discretizations.
- In order to avoid discrepancies in how boundary integrals are treated, the solution should satisfy homogeneous Neumann conditions on the non-Dirichlet portion of the boundary (zero surface traction).
- The domain must be non-trivial, in the sense that a coarse polyhedral FEM discretization can not exactly capture its shape.
- Optimal convergence rates are not generally attained for real-world problems unless graded or adaptive meshes are employed. Since we would like to study convergence rates for regular meshes, the solution must be sufficiently regular as well (smooth and without sharp gradients).

Since we can not use MMS for the aforementioned reasons, and we are not aware of any published problems with exact solutions that would fit the requirements we have laid out, we resort to computing a high-resolution reference solution and compare a set of lower-resolution discretizations to the reference solution.

1.2 Problem description

We study the deformation of a linearly elastic material in a static equilibrium setting, with Poisson's ratio $\nu = 0$ and Young's modulus $E = 5 \cdot 10^6$ Pa. Consider the unit ball B_1 and the half-space $H \subseteq \mathbb{R}^3$ defined by $y \geq 0$. The intersection $B_1 \cap H$ is a hemisphere, on whose flat region ($y = 0$) we impose the Dirichlet condition $\mathbf{u} = 0$. The remaining part of the boundary is assumed to have zero traction. We thus consider the PDE in weak form

$$-\int_{\Omega} \mathbf{P}(\mathbf{F}) : \nabla \mathbf{w} \, d\mathbf{X} + \int_{\Omega} \mathbf{f}^{\text{ext}} \cdot \mathbf{w} \, d\mathbf{X} = 0 \quad \forall \mathbf{w} \in V.$$

Since we do not have an exact solution available, the choice of \mathbf{f}^{ext} is crucial to reproducing optimal convergence rates. A simple choice would be the standard gravitational force. However, this leads to localized stress concentrations near the structural weak points of the domain ($y = 0, x^2 + z^2 = 1$), which appears to preclude higher-order convergence rates with regular meshes, presumably due to a lack of sufficient regularity in the exact solution. Therefore we instead use an artificial force which is concentrated at the top of the hemisphere and fades out before reaching the aforementioned structural weak points. The resulting deformation is a slight dent




Fig. 1. The reference domain of the hemisphere experiment used for experimental convergence analysis.

on the top of the hemisphere. Specifically, we define the body force \mathbf{f}^{ext} as

$$\mathbf{f}^{\text{ext}}(x, y, z) = -\alpha \psi \left(3\sqrt{x^2 + z^2} \right) y \hat{\mathbf{y}},$$

where $\hat{\mathbf{y}} = (0, 1, 0)$, $\alpha := 5 \cdot 10^5$ and ψ is the standard bump function defined by

$$\psi(r) = \begin{cases} \exp\left(-\frac{1}{1-r^2}\right), & r \in (-1, 1) \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

Note that the composition $\psi(3\sqrt{x^2 + y^2})$ is smooth even though $\sqrt{x^2 + y^2}$ is not.

1.3 Discretization

In principle, our embedded method could handle curved boundaries if we were to use an appropriate subdivision integration method that can accurately deal with curved interfaces and couple it with our simplification algorithm. However, it would be difficult with our current setup to compute an accurate reference solution on such a curved domain, and our focus is in any case on polyhedral domains. Therefore we instead consider a tetrahedral approximation of the hemisphere with 10395 tetrahedra generated by Gmsh [Geuzaine and Remacle 2009], depicted in Figure 1. This means that the *exact* domain Ω is a polyhedron, not a real hemisphere. As a reference solution, we apply two rounds of uniform refinement to produce a boundary-conforming tetrahedral mesh with 665k elements and 120k vertices. We use cubic tetrahedral elements for the finite element discretization, so that the total node count is 3.1M nodes.

We consider a family of uniform hexadral meshes with $(2i) \times (2i) \times i$ cells, where $i \in \{1, 2, 4, 8, 16, 24, 32\}$. Each uniform mesh covers the domain $[-1, 1] \times [0, 1] \times [-1, 1]$. The mesh cell width h is then defined by $h := 1/i$. For each resolution i , all cells that lie completely outside the embedded domain Ω are discarded, and the remaining cells form the background mesh Ω_h^i . We consider elements Hex8 (trilinear hexahedra) and Hex20 (quadratic Serendipity elements). We solve the problem in two ways: by simulating on Ω_h^i with the standard FEM, and by embedding the exact domain Ω into Ω_h^i with the algorithm from Section 4 to form quadrature rules that we use

for the FCM. Since we focus on order of convergence for the finite element spaces, we would like to eliminate systematic errors due to inaccurate integration. In all experiments, we use sufficiently high-order quadratures so that quantities are exactly integrated (in the case of the stiffness matrix and internal forces) or to very high precision with polynomial order 10 in the case of the external force, which contains a non-polynomial exponential function. For the stiffness matrix and internal force, we used our simplification algorithm, which gave the same results as in the non-simplified case, as expected.

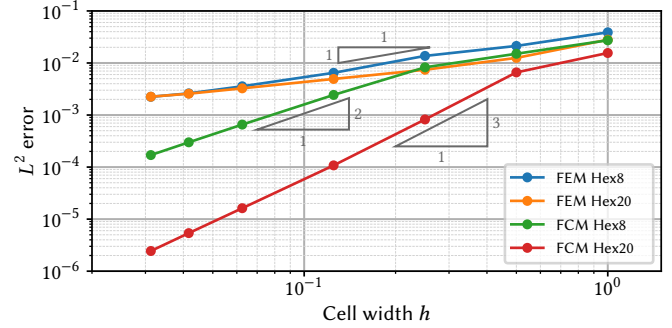


Fig. 2. L^2 errors relative to a high-resolution reference solution for a static equilibrium problem.

To measure the errors, we compute the squared L^2 error in each element of the *reference discretization* (cubic tetrahedra) by interpolating the displacements of the reference solution and the FEM/FCM approximation at quadrature point locations in the cells of the reference mesh. The resulting L^2 error as a function of mesh width h is depicted in Figure 2. The expected convergence rates for a Finite Element space with polynomial degree p is $O(h^{p+1})$ [Brenner and Scott 2007], which corresponds to h^2 for Hex8 and h^3 for Hex20. Here we see how the standard FEM (which does not treat the embedded domain accurately) experiences a severe reduction in convergence rate. In this case, the overall error is dominated by the local error near the hemisphere surface, which is a result of the poorly captured geometry. Moreover, the higher-order Hex20 elements are unable to improve over the (already poor) Hex8 results, suggesting that the main benefits of a higher-order discretization are essentially lost when the geometry is poorly approximated. In contrast, the FCM results show optimal convergence rates, though the rate for Hex20 diminishes somewhat for higher resolutions (lower h). Since an exact solution is not available, we are not able to ascertain with certainty the cause, but some plausible factors that come into play are:

- The (numerical) reference solution is perhaps not accurate enough.
- The exact solution may not be sufficiently regular.
- The handling of Dirichlet boundary conditions is not optimal, in the sense that some nodes of the background mesh at $y = 0$ are constrained to have zero displacement even though they lie outside of the embedded domain Ω . This might *perhaps* overconstrain the solution space for higher-order elements.

Nevertheless, the results demonstrate that the FCM coupled with our embedded quadrature algorithm is able to attain higher order convergence rates, and clearly outperforms the standard FEM for hexahedral meshes where the geometry can not be accurately captured. This is especially true for higher-order elements, such as the quadratic Serendipity Hex20 elements.

REFERENCES

- Susanne Brenner and Ridgway Scott. 2007. *The mathematical theory of finite element methods*. Vol. 15. Springer Science & Business Media.
- Christophe Geuzaine and Jean-François Remacle. 2009. Gmsh: A 3-D finite element mesh generator with built-in pre-and post-processing facilities. *International journal for numerical methods in engineering* 79, 11 (2009), 1309–1331.
- Patrick J Roache. 2002. Code verification by the method of manufactured solutions. *Journal of Fluids Engineering* 124, 1 (2002), 4–10.