

Coupling volume-of-fluid based interface reconstructions with the extended finite element method

John Dolbow^{a,*}, Stewart Mosso^b, Joshua Robbins^b, Tom Voth^b

^a *Department of Civil and Environmental Engineering, Duke University, Durham, NC 27708-0287, USA*

^b *Sandia National Laboratories, Albuquerque, NM 87185-0378, USA¹*

Received 3 March 2007; received in revised form 2 July 2007; accepted 4 August 2007

Available online 7 September 2007

Abstract

We examine the coupling of the patterned-interface-reconstruction (PIR) algorithm with the extended finite element method (X-FEM) for general multi-material problems over structured and unstructured meshes. The coupled method offers the advantages of allowing for local, element-based reconstructions of the interface, and facilitates the imposition of discrete conservation laws. Of particular note is the use of an interface representation that is volume-of-fluid based, giving rise to a segmented interface representation that is not continuous across element boundaries. In conjunction with such a representation, we employ enrichment with the ridge function for treating material interfaces and an analog to Heaviside enrichment for treating free surfaces. We examine a series of benchmark problems that quantify the convergence aspects of the coupled method and examine the sensitivity to noise in the interface reconstruction. The fidelity of a remapping strategy is also examined for a moving interface problem.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Interface reconstruction; Extended finite element; Volume-of-fluid

1. Introduction

The extended finite element method (X-FEM) introduced in Moës et al. [12], Dolbow et al. [5], and Daux et al. [4] was designed to facilitate the simulation of evolving discontinuities and interfaces without continuous remeshing. It does so by building appropriately selected enrichment functions into a standard finite-element basis through the partition-of-unity framework [9]. Implicit in this method is a geometric representation of the feature of interest, be it a crack surface or phase interface, for example, that may be independent of the finite element mesh. Existing approaches include classical and level-set

representations [15]. In this work, we examine the advantages of employing volume-of-fluid interface representations in conjunction with the X-FEM.

While originally designed to facilitate the simulation of crack propagation in brittle materials, the X-FEM has emerged as an attractive candidate for evolving interface problems. The method has been applied to simulate stable [8] and unstable [6,20] phase interface evolution, two-phase fluid flow [3], and resin-transfer molding [10], just to name a few. Of course there exist many alternative methods for all of these problems; however, it does appear that the accuracy of the X-FEM is superior to many. For example, in a side-by-side comparison, Vaughan et al. [17] recently showed the X-FEM to be more accurate than the immersed interface method for elliptic interface problems.

Most multi-material Eulerian or Arbitrary Lagrangian–Eulerian methods allow for single elements or cells in the computational domain to contain volume fractions of various constituents [1]. These same methods typically employ volume-of-fluid type interface reconstruction schemes over

* Corresponding author.

E-mail address: jdolbow@duke.edu (J. Dolbow).

¹ Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Security Administration under Contract DE-AC04-94AL85000.

such “mixed elements” in order to improve flux calculations. It is widely recognized that a drawback to multi-material finite element methods concerns the smearing of discontinuities across the interface. Recently, Vitali and Benson [18] incorporated ideas from the X-FEM into a multi-material code to improve the contact kinematics in mixed elements. They employed structured Cartesian meshes and Youngs’ algorithm to reconstruct the interface. The spatial convergence of the method and issues associated with the segmented representation of the interface were not examined.

In this paper, we employ the second-order patterned-interface-reconstruction (PIR) algorithm proposed by Mosso et al. [13] for arbitrary, unstructured meshes. The PIR method was developed as a second-order accurate successor to Youngs’ method [19]. Both Youngs’ method and PIR are used within volume-of-fluid (VOF) formulations [2] to accurately partition the volume fluxed (advected) from multi-material (mixed) donor elements to neighboring elements. An advantage of the VOF formulation is volume conservation during the advection step. In particular, the interface is constructed in such a way that material volumes are strictly conserved during fluxing. The “classical” level-set method does not, in general, give rise to volume conservation. In fact, one approach to address this issue is to couple level sets to volume-of-fluid reconstructions, as in Sussman [16].

Youngs’ method (developed for Cartesian, structured grids) uses piecewise linear interfaces to partition mixed-material elements. Youngs’ algorithm consists of (i) interface normal approximation and (ii) volume conserving interface positioning. The interface normal is first determined using volume fraction gradients as calculated from the immediate mixed-material element neighborhood. The interface is then positioned, using the normal of step (i) to conserve donor element volumes. Inaccuracies in the algorithm’s gradient approximated normal result in an algorithm that cannot, in general, exactly reproduce a planar material interface and will not produce a second-order representation of most curved interfaces.

In Mosso et al. [14], a two-dimensional smoothing algorithm was developed that reproduced linear interfaces in all orientations and relative element positions. The algorithm was developed on non-Cartesian, unstructured grids with arbitrary elements. The method incorporates steps (i) and (ii) above and includes the concept of stability points to improve interface normal accuracy. In two-dimensions, stability points are the midpoints of the reconstructed interfaces. By locally fitting a plane to the stability points (i.e. “smoothing” the interface), the interface normal is improved relative to step (i). By iterating over the normal improvement and positioning steps, linear interfaces are reproduced within a few iterations. A drawback to this algorithm is that it only utilizes two neighboring stability points chosen from among the larger number of neighboring stability points. For curved interfaces, the choice of these neighbors was critical to the iterative improvement

of the local interface normal. The PIR method [13] employed here uses a larger number of stability points in its smoothing algorithms, and has been shown to obtain second-order accuracy on planar and curved interfaces.

Concerning the coupling to the X-FEM, there are many differences between the current approach to interface representation and existing classical or level-set representations. Classical representations typically represent the interface as a series of connected segments or facets. With the VOF-based method proposed herein, by contrast, the interface is represented by a single plane within each element. This is also in contrast with level-set representations, where the interface is usually reconstructed on the basis of nodal (or vertex) signed distance fields. With linear triangular or tetrahedral elements, no subdivision of the elements is required for integration with the VOF representation. Further, the element-level volume fractions are easily used to determine which nodes are selected for enrichment.

This paper is organized as follows. In Section 2, we provide the governing equation and weak form for simple one-sided and two-sided problems that will be used to investigate the coupled method. In Section 3, we describe the numerical discretization with the coupled X-FEM/VOF method, including details of the interface reconstruction algorithm. The fidelity of the coupled method is then investigated with a series of benchmark interface problems in Section 4. Finally, we provide a summary and concluding remarks in the last section.

2. Formulation

In this section, we provide the governing equations and equivalent weak formulations for two fundamental problems that will be used to explore the coupled X-FEM/VOF method.

2.1. One-sided problem

Consider the “one-sided problem” described by an interface \mathcal{S} partitioning the domain \mathcal{R} into the disjoint sets \mathcal{R}^+ and \mathcal{R}^- , as shown in Fig. 1.

We consider a simple, one-sided Poisson problem of the form

$$\Delta u = f \quad \text{in } \mathcal{R}^-, \quad (2.1a)$$

$$u = \mathbf{u}_d \quad \text{on } \partial\mathcal{R}_d, \quad (2.1b)$$

$$\nabla u \cdot \mathbf{n}_o = \mathbf{m} \quad \text{on } \partial\mathcal{R}_n, \quad (2.1c)$$

$$\nabla u \cdot \mathbf{n} = \mathbf{g} \quad \text{on } \mathcal{S}, \quad (2.1d)$$

where Δ is the standard Laplace operator, and \mathbf{n}_o is the outward unit normal to $\partial\mathcal{R}_n$, as shown in Fig. 1. We assume that the boundary $\partial\mathcal{R}^-$ is composed of the disjoint sets of Dirichlet $\partial\mathcal{R}_d$, Neumann $\partial\mathcal{R}_n$ and interface \mathcal{S} parts such that $\overline{\mathcal{S} \cup \partial\mathcal{R}_d \cup \partial\mathcal{R}_n} = \overline{\partial\mathcal{R}^-}$. The above is a simplification of classical one-sided Stefan problems that arise from models for a wide range of evolving interface phenomena.

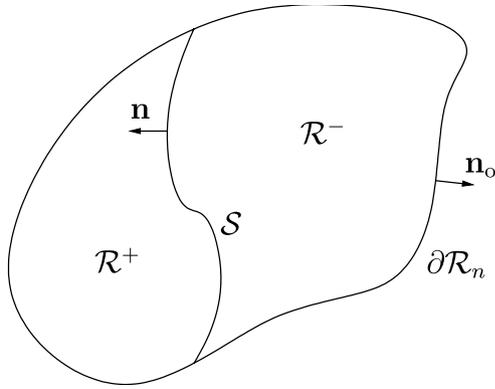


Fig. 1. Notation for the one-sided problem. A domain \mathcal{R} partitioned into regions \mathcal{R}^+ and \mathcal{R}^- by the interface \mathcal{S} . The normal \mathbf{n} to the interface is defined such that it points outward from the \mathcal{R}^- subdomain.

2.1.1. Standard weak formulation

We write \mathcal{U} for the space of admissible bulk fields, and \mathcal{V} for the corresponding space of variations. The weak form reads: Find $u \in \mathcal{U}$ such that

$$\int_{\mathcal{R}^-} \nabla \mathbf{w} \cdot \nabla \mathbf{u} dv = \int_{\mathcal{S}} \mathbf{w} \mathbf{g} da + \int_{\partial\mathcal{R}_n} \mathbf{w} \mathbf{m} da \quad (2.2)$$

for all $w \in \mathcal{V}$.

2.2. Two-sided problem

In the “two-sided problem” the governing equations of static equilibrium are enforced in all of the domain \mathcal{R} so that

$$\text{Div } \boldsymbol{\sigma} = 0 \quad \text{in } \mathcal{R}, \quad (2.3a)$$

$$\mathbf{u} = \mathbf{u}_d \quad \text{on } \partial\mathcal{R}_d, \quad (2.3b)$$

$$\boldsymbol{\sigma} \mathbf{n}_o = \mathbf{t} \quad \text{on } \partial\mathcal{R}_n, \quad (2.3c)$$

where $\boldsymbol{\sigma}$ is the stress tensor. The strain $\boldsymbol{\epsilon}$ and isotropic elastic material response are given by

$$\boldsymbol{\epsilon} = \nabla_s \mathbf{u} = \frac{1}{2} (\nabla \mathbf{u} + \mathbf{u} \nabla), \quad (2.4a)$$

$$\boldsymbol{\sigma} = C^+ \boldsymbol{\epsilon} \quad \text{in } \mathcal{R}^+, \quad \boldsymbol{\sigma} = C^- \boldsymbol{\epsilon} \quad \text{in } \mathcal{R}^-, \quad (2.4b)$$

and C is the standard isotropic elastic stiffness tensor. On the interface, continuity of tractions and displacements are enforced, i.e.

$$[[\mathbf{u}]] = \mathbf{0} \quad \text{on } \mathcal{S}, \quad (2.5a)$$

$$[[\boldsymbol{\sigma} \mathbf{n}]] = \mathbf{0} \quad \text{on } \mathcal{S}, \quad (2.5b)$$

where $[[\cdot]]$ is the jump operator.

2.2.1. Standard weak formulation

The weak form for (2.3c) is given by: find $\mathbf{u} \in \mathcal{U}$ such that

$$\int_{\mathcal{R}^+} \nabla_s \mathbf{w} : C^+ : \nabla_s \mathbf{u} dv + \int_{\mathcal{R}^-} \nabla_s \mathbf{w} : C^- : \nabla_s \mathbf{u} = \int_{\partial\mathcal{R}_n} \mathbf{w} \cdot \mathbf{t} da \quad (2.6)$$

for all $\mathbf{w} \in \mathcal{V}$. The above weakly enforces traction continuity across the interface \mathcal{S} .

3. Numerical discretization

We consider a quasi-uniform partition \mathcal{R}^h of the domain \mathcal{R} into non-overlapping element domains \mathcal{R}^e with boundaries $\partial\mathcal{R}^e$. We assume the volume fractions f_e^+ and f_e^- of the positive and negative domains

$$f_e^+ = \frac{\text{meas}(\mathcal{R}^e \cap \mathcal{R}^+)}{\text{meas}(\mathcal{R}^e)}, \quad f_e^- = \frac{\text{meas}(\mathcal{R}^e \cap \mathcal{R}^-)}{\text{meas}(\mathcal{R}^e)} \quad (3.1)$$

are given for each element. Since we consider only positive and negative regions, the two quantities are obviously related by

$$1 = f_e^+ + f_e^-. \quad (3.2)$$

An element volume fraction in the range $0 < f_e^+ < 1$ indicates that the interface is located within the element. Accordingly, we represent the interface in the element by a plane that partitions the element domain into portions that exactly correspond with the volume fractions. The plane can be associated with an element normal \mathbf{n}_e . We note that this plane is not unique, and discuss this interface reconstruction based on volume fractions in Section 3.3. It should be clear at this point, however, that the interface planes will not generally align with element boundaries, and further may not be contiguous between adjacent elements. We write \mathcal{S}^h for the set of all element interface planes making up the discrete representation of the interface. For the sake of concreteness, an example of \mathcal{R}^h and \mathcal{S}^h is shown in Fig. 2.

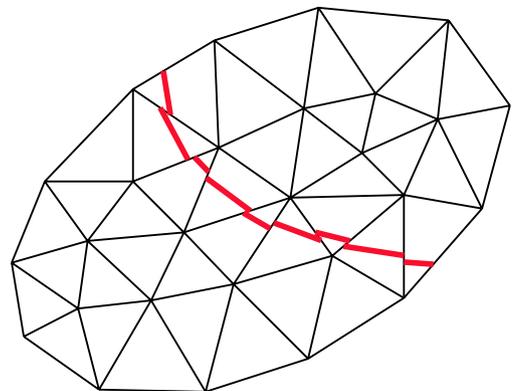


Fig. 2. Two-dimensional example of a bulk partition of the domain into a finite element mesh and corresponding partition of the interface into a set of element planes.

3.1. One-sided problem

We now consider an approximation for the field u and weight function w of the form

$$u^h(\mathbf{x}) = \sum_{i \in I} N_i(\mathbf{x}) u_i, \quad (3.3a)$$

$$w^h(\mathbf{x}) = \sum_{i \in I} N_i(\mathbf{x}) w_i, \quad (3.3b)$$

where N_i are the nodal shape functions, and I denotes the set of nodes

$$I = \{j | \bar{\omega}_j \cap \mathcal{R}^- \neq \emptyset\}, \quad (3.4)$$

which have some portion of their supports ω_j (with closure $\bar{\omega}_j$) intersecting the domain \mathcal{R}^- . This is necessary to maintain a partition-of-unity over \mathcal{R}^- which will allow for good convergence in the error $\|u^h - u\|$.

To determine the nodes whose basis functions contribute to the approximation in the one-sided problem, we take advantage of the element volume fraction information to calculate nodal-support volume fractions. For example, by looping over the set of elements connected to a vertex, we can easily calculate

$$f_n^- = \frac{\text{meas}(\omega_n \cap \mathcal{R}^-)}{\text{meas}(\omega_n)}, \quad (3.5)$$

the nodal-support volume fraction (on the negative side) of node n . We exclude from the set I those nodes whose volume fractions are below a user-specified tolerance. The approach is analogous to that used with the extended finite element method with classical or level-set interface representations.² The one-sided problem has gained interest as a basic problem to study the X-FEM of late, as two one-sided problems (with overlapping elements near the interface) is equivalent to enrichment with a Heaviside function.

By substituting (3.3) into a Galerkin approximation to the weak form (2.2), we arrive at the system of equations

$$\mathbf{Kd} = \mathbf{f}, \quad (3.6)$$

where the stiffness matrix \mathbf{K} and forcing vector \mathbf{f} are assembled from element vectors, i.e.

$$\mathbf{K} = \bigwedge \mathbf{k}_e, \quad \mathbf{f} = \bigwedge \mathbf{f}_e. \quad (3.7)$$

We note that in the case of linear triangle and tetrahedral elements, the element stiffness matrices can be obtained by simply scaling the classical stiffness matrix by the element volume fraction f_e^- . In this case, it is not necessary to partition any element into “subelements” (as is customary in the X-FEM) for integration purposes. This greatly simplifies and expedites the computation of the element stiffness matrices.

3.2. Two-sided problem

For the two-sided problem, the field \mathbf{u} and weight functions \mathbf{w} are approximated as

$$\mathbf{u}^h(\mathbf{x}) = \sum_{i \in I} N_i(\mathbf{x}) \mathbf{u}_i + \sum_{j \in J} \psi^e(\mathbf{x}) N_j(\mathbf{x}) \mathbf{a}_j, \quad (3.8a)$$

$$\mathbf{w}^h(\mathbf{x}) = \sum_{i \in I} N_i(\mathbf{x}) \mathbf{w}_i + \sum_{j \in J} \psi^e(\mathbf{x}) N_j(\mathbf{x}) \mathbf{b}_j, \quad (3.8b)$$

where \mathbf{a}_j denotes the additional degrees of freedom for enrichment, J is the set of nodes

$$J = \{j | \bar{\omega}_j \cap \mathcal{S} \neq \emptyset\}, \quad (3.9)$$

and $\psi^e(\mathbf{x})$ is the ridge enrichment function proposed by Moës et al. [11]

$$\psi^e(\mathbf{x}) = \sum_{i \in I} |\phi_i^e| N_i(\mathbf{x}) - \left| \sum_{i \in I} \phi_i^e N_i(\mathbf{x}) \right|. \quad (3.10)$$

The reconstructed interface is included in the nodal enrichment through ϕ_i^e , which are the signed normal distances from the local intersecting plane to the nodes composing the element. One can regard ϕ_i^e as a level-set representation of the reconstructed interface that is defined only locally on element e . We note that, due to the segmented nature of the interface representation, ϕ_i^e will not necessarily be the same value for the same node in adjacent elements.

3.3. Volume fraction based interface reconstruction

The PIR method employed here involves two key steps: (i) an approximation to the interface normal in each element; and (ii) volume preserving interface repositioning. Both steps are facilitated by the use of smoothing algorithms. In this section, we summarize the main aspects of the method. Interested readers are referred to Mosso et al. [13] for details.

The reconstruction begins by partitioning the set of elements in the mesh into a list of mixed elements (with multiple materials, and thus at least one interface) and a list of all single material elements. The first step is to approximate the interfacial normal in each mixed element, by performing a gradient calculation. In particular, the unit normal \mathbf{n} to the interface is approximated in each mixed element using a Dukowicz [7] surface integral formulation as

$$\mathbf{n}_e \simeq \frac{-\nabla f}{|\nabla f|} = \frac{-\oint f \hat{\mathbf{s}} dS}{|\oint f \hat{\mathbf{s}} dS|}, \quad (3.11)$$

where f is a discrete volume fraction field, and $\hat{\mathbf{s}}$ is the unit outward normal to the control volume. Vertex values for the field f are determined by considering the element volume fractions of the material in those elements connected to the vertex, as illustrated in Fig. 3. The integrals in (3.11) (using constant f) are formed over the surface of each vertex control volume.

Once the normal is calculated with (3.11), the interface is positioned in each mixed element using the element volume

² For additional insight and details, see Daux et al. [4] and Sukumar et al. [15].

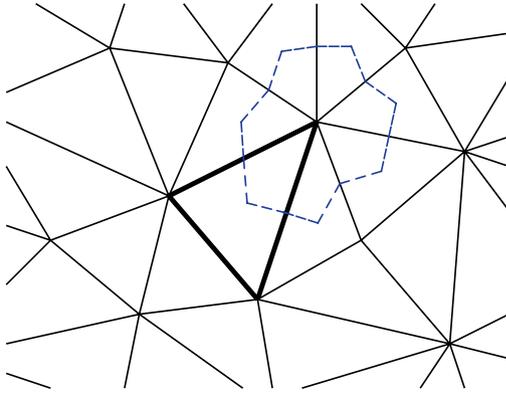


Fig. 3. Illustration of volume fraction gradient computation. The algorithm will calculate the gradient in the center element (indicated with heavy lines). The volume fractions will be averaged for each vertex using control volumes. The dashed lines indicate the boundary of the control volume for the top-most vertex. The edges of the control volume connect element centroids-to-centroids of element edges.

fraction f_c . This is repeated for each mixed element in the mesh, giving rise to an (initial) interface reconstruction. The PIR method then employs planar and spherical smoothing algorithms to iteratively improve upon this reconstruction. Both smoothing algorithms rely on the stability points.

For the first iteration, only the planar smoothing is employed because the initial normal given by (3.11) can be noisy for unstructured grids. Consider a linear interface passing through a neighborhood of the mesh as shown in Fig. 4. Even with element volume fractions specified to match a linear interface, the gradient normals given by (3.11) will yield reconstructed interfaces in each element

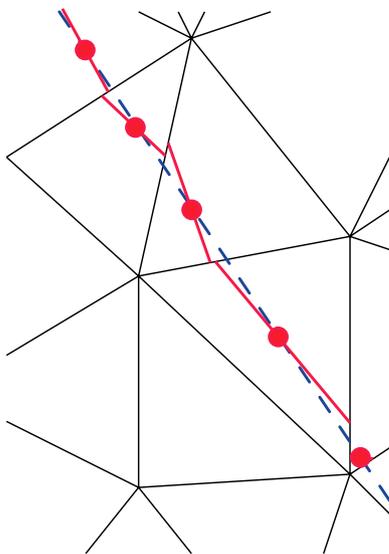


Fig. 4. Illustration of input configuration to planar smoothing using stability points. The dotted line is the given straight interface, and the solid lines in each element indicate the initial interfacial segments. The filled circles are the stability points, located at the centroid of each segment.

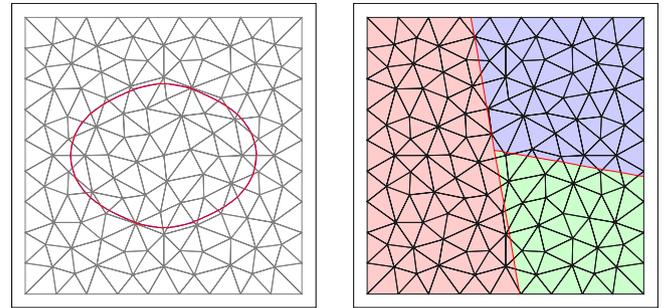


Fig. 5. Sample PIR reconstructed interfaces. Two-dimensional curved non-spherical test on an unstructured grid (left), and triple point, planar interfaces on an unstructured grid (right).

that do not align with the exact interface. Stability points are located at the centroid of each reconstructed interface as shown in Fig. 4. The stability points represent a good approximation for the position of the interface if the normal were more accurately chosen. The planar smoothing algorithm performs a least-squares fit of a plane to the appropriate, neighboring stability points about each mixed element.

The second smoothing method is similar, and is based on the assumption of a curved interface. It computes a center of convergence based upon a circular chord extending from the element’s stability point to each of the filtered, neighboring stability points. By repositioning the planar interface in the mixed element such that the new interface is orthogonal to the radial extending from the center of convergence to the element’s new stability point, second-order accuracy is achieved for curved interfaces.

For each material, both smoothing algorithms are initially used. To select the smoothing result for each element, the interface is extended from the home element to its neighboring interfaces and the volume between the extrapolated interface and the interface in the neighboring elements is computed. The smoothing algorithm that produced the lowest volumetric “discrepancy” or “roughness” is used in further smoothing iterations for that material interface fluxing step. The algorithm is fast and typically converges in two to four iterations. The robustness of the algorithm has been demonstrated by its performance on a series of benchmark tests. Sample PIR reconstructed interfaces are shown for a curved interface and triple point, planar interfaces in Fig. 5.

4. Numerical examples

We first consider two-dimensional problems where the interface is fixed. We consider a scalar one-sided problem as well as a displacement-based material inclusion problem. These problems possess analytical solutions, allowing us to examine convergence and the role of perturbations to the interface geometry. We then provide results for an artificial evolving interface problem to illustrate the fidelity of the remap/reconstruction algorithms.

With the one-sided problem, a Neumann condition is enforced on the interface, and the nodal selection procedures of the X-FEM are employed near the interface. To a degree, this is analogous to enrichment with the Heaviside function. The problem therefore tests the sensitivity of the accuracy in the flux imposition to the interface geometry. With the two-sided problem, displacement continuity is enforced on the interface through the use of the ridge enrichment function. This problem therefore tests the sensitivity of such a constraint to the interface geometry, as represented by a set of disjoint segments. Finally, the evolving interface problem couples the X-FEM solution to the remapping and reconstruction algorithms repeatedly, and examines the extent to which errors propagate.

4.1. The stationary one-sided problem

We consider the two-dimensional problem of a circular void with boundary $\mathcal{S} = \{\mathbf{x} : |\mathbf{x}| = R_s\}$ embedded in a rectangular domain as shown in Fig. 6. We consider the function

$$u(r) = c_1 + c_2 \log r, \tag{4.1}$$

in terms of radial coordinate r , measured from the center of the circular void. It can be shown that, regardless of the specified constants c_1 and c_2 , such a function satisfies the Poisson Eq. (2.1a) with $f=0$. Accordingly, we impose a strict Dirichlet condition on the outer boundary $\partial\mathcal{R}_d = \partial\mathcal{R}^-$ consistent with (4.1), and prescribe the flux $\mathbf{g} = -c_2/R_s$ on \mathcal{S} .

We consider a sequence of meshes of increasing refinement, and use h to denote the characteristic mesh size. The coarsest mesh and segmented interface are shown in Fig. 7. A close-up view of the interface near the void is also shown in Fig. 7. It can be seen that the interface geometry is not continuous, and further that it does not conform to the underlying mesh. In this case, element volume fractions were initialized according to the exact intersection of the element with the circular interface, and the interface was then reconstructed using the PIR algorithm described in Section 3.3.

Fig. 8 shows the X-FEM/VOF solution over an intermediate mesh, and Fig. 9 shows the results of the convergence

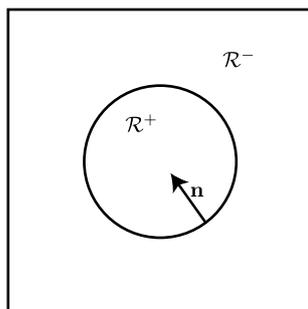


Fig. 6. Geometry and notation for one-sided benchmark problem.

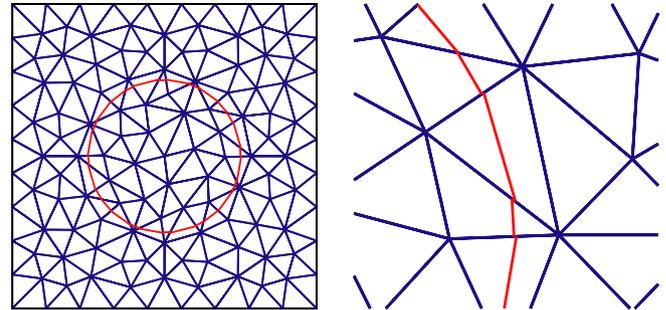


Fig. 7. Coarse mesh (left) and zoom of interface (right) for one-sided problem.

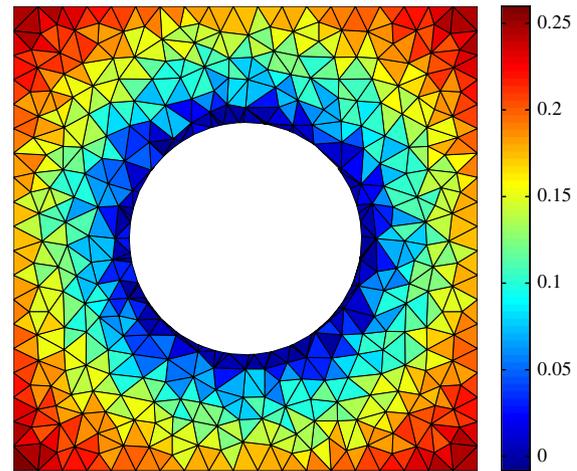


Fig. 8. Numerical approximation u^h to the solution u of the one-sided problem.

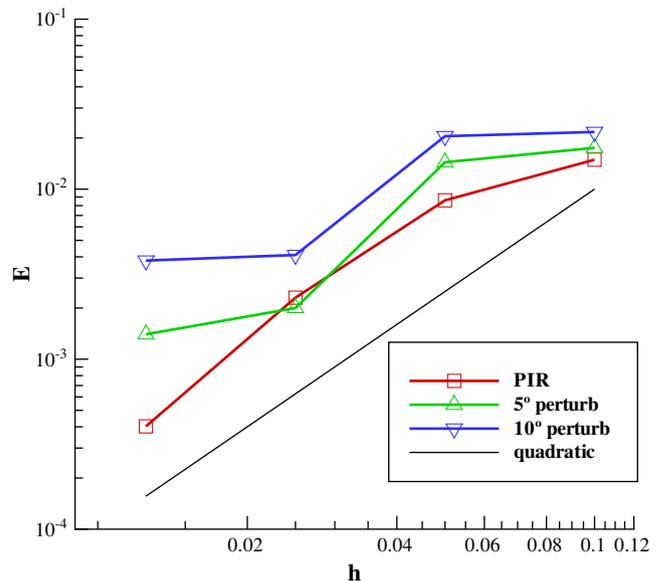


Fig. 9. Convergence in L_2 error for the one-sided problem.

study for this problem. Here, we calculate the norm of the error in L_2 ,

$$E = \frac{\left(\int (u - u^h)^2 dv\right)^{1/2}}{\left(\int u^2 dv\right)^{1/2}}. \tag{4.2}$$

With the interface reconstruction algorithm described in Section 3.3, we report an optimal rate of convergence. An examination of the error in the interface geometry reveals that it is also second-order. Conversely, if we introduce a random perturbation to the interface geometry, with each segment perturbed up to 5° or 10°, the convergence is clearly affected as shown in Fig. 9. This is to be expected, as the perturbation to the interface destroys the second-order accurate representation for the geometry. These studies provide testament to the need to employ accurate interface reconstruction algorithms in conjunction with the X-FEM.

4.2. The stationary two-sided problem

Convergence behavior of the two-sided problem formulation is examined in the context of a circular inclusion in a circular matrix as shown in Fig. 10. The inclusion is given a non-zero initial strain and then allowed to equilibrate with the surrounding material. The appropriate Dirichlet boundary conditions are applied so that the computation can be performed over a single quadrant of the actual problem.

For this problem, we first examine results using structured meshes of four-node quadrilateral elements. The coarsest mesh used for the convergence study is shown in Fig. 11 with the reconstructed interface superposed. It is clear from the figure that the interface is discontinuous and does not conform to the mesh. To treat the void space, we employ a nodal selection strategy analogous to that described for the one-sided problem.

Convergence is measured in terms of the energy norm

$$e = \frac{\left(\int_{\mathcal{R}} \nabla_s(\mathbf{u}^h - \mathbf{u}) : C : \nabla_s(\mathbf{u}^h - \mathbf{u})\right)^{1/2}}{\left(\int_{\mathcal{R}} \nabla_s \mathbf{u} : C : \nabla_s \mathbf{u}\right)^{1/2}}. \tag{4.3}$$

Fig. 12 shows optimal (first order) convergence in the energy norm when using the interface reconstruction of Section 3.3. As in the one-sided problem, random perturbations in the interface geometry produce a significant

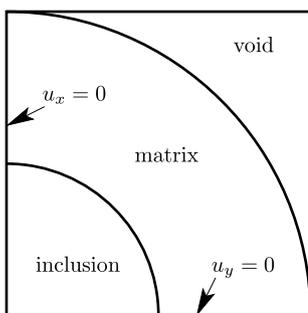


Fig. 10. Geometry and notation for two-sided benchmark problem.

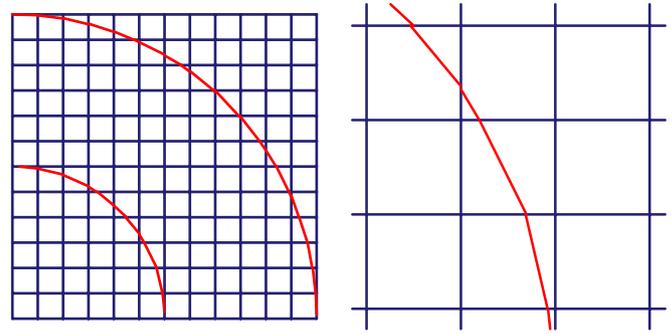


Fig. 11. Coarse structured mesh and zoom near interface for two-sided problem.

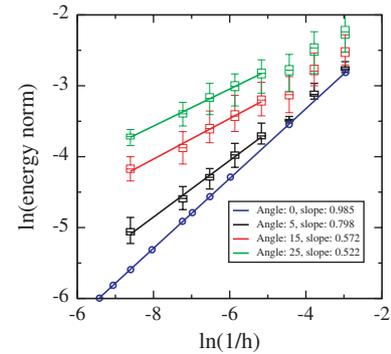


Fig. 12. Convergence in energy norm for the two-sided problem. Angles listed in the legend indicate the degree of random perturbation given to the interfacial normal. Error bars are indicative of the range of results obtained using various random seeds in each case.

drop in the rate of convergence. At each resolution and angle of perturbation, a series of calculations was run with unique random seeds to produce the distribution shown in the convergence plot as error bars.

Results using unstructured meshes of quadrilateral elements compare favorably. We conducted convergence studies using random perturbations of 5° and 15° to interface segments. The coarsest unstructured mesh is shown in Fig. 13. Nearly optimal rates of convergence are again observed when using the PIR algorithm. We report a small decrease in the rate of convergence and accuracy as compared to the results on the structured mesh, as expected.

4.3. Evolving interface with remapping

We now investigate the ability of the method to handle evolving interface problems, in which remapping volume fractions between two meshes is a key step. To effect this, we again develop approximate solutions to the one-sided problem with a circular interface. Using the reference mesh \mathcal{M}^r with nodal coordinates \mathbf{x}^r , we determine the coordinates \mathbf{x}^d of a deformed mesh \mathcal{M}^d with the simple algorithm

$$\mathbf{x}_I^d = \mathbf{x}_I^r + \Delta t(c_1 - u_I \mathbf{n}), \quad I = 1, \dots, N, \tag{4.4}$$

where Δt is a prescribed, fixed time step, and u_I is the nodal degree-of-freedom obtained by the X-FEM. As such, this

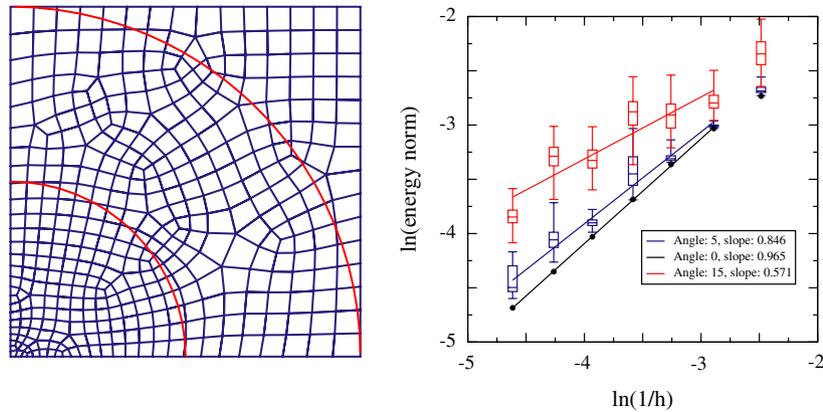


Fig. 13. Coarse unstructured mesh (left) and corresponding convergence results (right) for two-sided problem.

represents a coupled X-FEM/PIR evolving interface problem.

The volume fractions in each element are maintained in the deformed mesh \mathcal{M}^d . We then clear the volume fractions in the reference mesh, and remap the volume fractions from \mathcal{M}^d to \mathcal{M}^r and begin anew. The effect is to evolve the interface outward. While this evolving-interface problem does not possess an analytical solution, the results are nonetheless illustrative.

A simple polyhedral intersection algorithm can be used to remap the volume fractions by intersecting the end of cycle position of the elements (donor elements), with the starting cycle position of the element and the element's neighbors (acceptor elements). The simplicity of the intersection is due to the convexity of the triangular elements. If the donor element contains only a single material, the volume of the polyhedral intersection is accumulated in the material volume of the intersected acceptor element. If the donor element were mixed, the polyhedral intersection would be intersected with the donor element's reconstructed interface. The volume of each interface partition is accumulated in the acceptor element material volumes. The polyhedral intersection algorithm takes the acceptor element polygon and intersects it by the planes describing each of the donor element's edges.

Fig. 14 shows the results obtained over a coarse mesh. We show the mesh and interface at equally-spaced time steps of $5\Delta t$. The coarseness of the mesh is apparent in

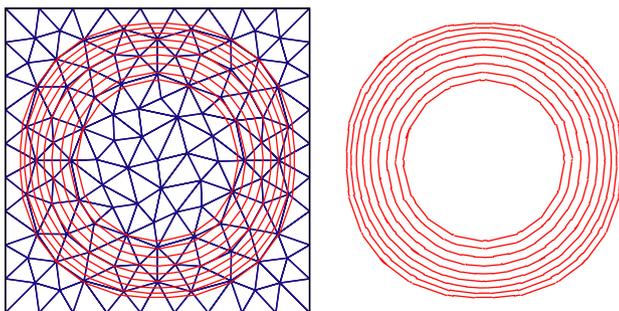


Fig. 14. Evolving interface problem: (left) mesh and interface at equally-spaced time steps; (right) interface segments only.

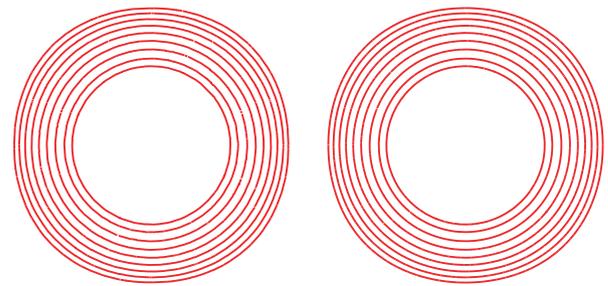


Fig. 15. Evolving interface geometry, at equally-spaced time steps on two subsequently refined meshes.

the results, with results over two subsequent refined meshes (Fig. 15) being much more smooth. In any case, the radial symmetry of the front is nicely preserved by the method.

5. Summary and concluding remarks

In this paper, we examined the use of the patterned-interface-reconstruction (PIR) algorithm in conjunction with the enrichment/nodal selection strategies of the extended finite element method (X-FEM). The method exhibits the advantages of local, element-based interface representation, in particular the ability to enforce strict volume conservation. We provided a series of numerical results for elliptic interface problems to illustrate the accuracy of the coupled method. The results clearly indicate that optimal rates of convergence are attainable with the coupled method, even though the interface is not strictly continuous between adjacent elements. The main contribution of this work over earlier efforts is seen to be the applicability to arbitrary, unstructured meshes, as well as the investigation of the convergence properties.

There exist a number of interesting extensions to this work that appear worth pursuing. If a continuous interface is deemed necessary, this method could certainly be further coupled to level-set representations to improve volume conservation, as has been demonstrated in the finite-difference literature [16]. These concepts might also be applied to the representation of fracture/failure processes, where

maintaining strict continuity of the evolving crack front can be challenging. Finally, we mention the extension to multi-material problems where triple junctions can easily occur in elements. Such problems are not easily treated with level-set representations.

Acknowledgements

John Dolbow would like to recognize the partial support of Sandia National Laboratories and the Department of Energy, to Duke University. The helpful comments of several reviewers at Sandia National Laboratories are also acknowledged.

References

- [1] D.J. Benson, Computational methods in Lagrangian and Eulerian hydrocodes, *Comput. Method Appl. Mech. Engrg.* 99 (1992) 235–394.
- [2] C.W. Hirt, B.D. Nichols, Volume-of-fluid (VOF) method for the dynamics of free boundaries, *J. Comput. Phys.* 39 (1981) 201–225.
- [3] J. Chessa, T. Belytschko, An extended finite element method for two-phase fluids, *J. Appl. Mech.* 70 (1) (2003) 10–17.
- [4] C. Daux, N. Moës, J. Dolbow, N. Sukumar, T. Belytschko, Arbitrary branched and intersecting cracks with the extended finite element method, *Int. J. Numer. Method Engrg.* 48 (2000) 1741–1760.
- [5] J. Dolbow, N. Moës, T. Belytschko, Discontinuous enrichment in finite elements with a partition of unity method, *Finite Element Anal. Des.* 36 (2000) 235–260.
- [6] J.E. Dolbow, E. Fried, H. Ji, A numerical strategy for investigating the kinetics of stimulus-responsive hydrogels, *Comput. Method Appl. Mech. Engrg.* 194 (2005) 4447–4480.
- [7] J.K. Dukowicz, Conservative rezoning (remapping) for general quadrilateral meshes, *J. Comput. Phys.* 54 (1984) 411–424.
- [8] H. Ji, D. Chopp, J.E. Dolbow, A hybrid extended finite element-level set method for modeling phase transformations, *Int. J. Numer. Method Engrg.* 54 (8) (2002) 1209–1233.
- [9] J.M. Melenk, I. Babuška, The partition of unity finite element method: basic theory and applications, *Comput. Method Appl. Mech. Engrg.* 139 (1996) 289–314.
- [10] N. Moës, E. Bechet, M. Tourbier, Imposing Dirichlet boundary conditions in the extended finite element method, *Int. J. Numer. Method Engrg.* 67 (2006) 1641–1669.
- [11] N. Moës, M. Cloirec, P. Cartraud, J.F. Remacle, A computational approach to handle complex microstructure geometries, *Comput. Method Appl. Mech. Engrg.* 192 (2003) 3163–3177.
- [12] N. Moës, J. Dolbow, T. Belytschko, A finite element method for crack growth without remeshing, *Int. J. Numer. Method Engrg.* 46 (1999) 131–150.
- [13] S. Mosso, C. Garasi, R. Drake, A smoothed two- and three-dimensional interface reconstruction method, *Comput. Visual. Sci.* (accepted for publication).
- [14] S. Mosso, B.K. Swartz, D.B. Kothe, A parallel, volume-tracking algorithm for unstructured meshes, in: P. Shiano, A. Ecer, J. Periaux, N. Satofuka (Eds.), *Parallel Computational Fluid Mechanics*, Elsevier, 1997, pp. 368–375.
- [15] N. Sukumar, D. Chopp, N. Moës, T. Belytschko, Modeling holes and inclusions by level sets in the extended finite element method, *Comput. Method Appl. Mech. Engrg.* 190 (2001) 6183–6200.
- [16] M. Sussman, A second order coupled level set and volume-of-fluid method for computing growth and collapse of vapor bubbles, *J. Comput. Phys.* 187 (2003) 110–136.
- [17] B.L. Vaughan, B.G. Smith, D.L. Chopp, A comparison of the extended finite element method with the immersed interface method for elliptic equations with discontinuous coefficients and singular sources, *Commun. Appl. Math. Comput. Sci.* 1 (1) (2006) 207–228.
- [18] E. Vitali, D.J. Benson, An extended finite element formulation for contact in multi-material arbitrary Lagrangian–Eulerian calculations, *Int. J. Numer. Method Engrg.* 67 (2006) 1420–1444.
- [19] D.L. Youngs, Time-dependent multi-material flow with large fluid distortion, in: K.W. Morton, J.J. Baines (Eds.), *Numerical Methods for Fluid Dynamics*, Academic Press, 1982, pp. 273–285.
- [20] N. Zabaras, B. Ganapathysubramanian, L.J. Tan, Modelling dendritic solidification with melt convection using the extended finite element method, *J. Comput. Phys.* 218 (1) (2006) 200–227.