INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS
Int. J. Numer. Meth. Fluids 2010; **63**:725–742
Published online 19 June 2009 in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/fld.2095

A conformal decomposition finite element method for modeling stationary fluid interface problems

David R. Noble*,†, Elijah P. Newren and Jeremy B. Lechman

Sandia National Laboratories Albuquerque, MS-0382, P.O. Box 5800 Albuquerque, New Mexico 87185-0382, U.S.A.

SUMMARY

A method is developed for modeling fluid transport in domains that do not conform to the finite element mesh. One or more level set functions are used to describe the fluid domain. A background, non-conformal mesh is decomposed into elements that conform to the level set interfaces. Enrichment takes place by adding nodes that lie on the interfaces. Unlike other enriched finite element methods, the proposed technique requires no changes to the underlying element assembly, element interpolation, or element quadrature. The complexity is entirely contained within the element decomposition routines. It is argued that the accuracy of the method is no less than that for eXtended Finite Element Methods (XFEM) with Heaviside enrichment. The accuracy is demonstrated using multiple numerical tests. In all cases, optimal rates of convergence are obtained for both volume and surface quantities. Jacobi preconditioning is shown to remove the ill-conditioning that may result from the nearly degenerate conformal elements. Copyright © 2009 John Wiley & Sons, Ltd.

Received 25 November 2008; Revised 16 April 2009; Accepted 18 April 2009

KEY WORDS: level set; enriched finite element; conformal decomposition; finite element; potential flow; multiphase flow

1. INTRODUCTION

Many fluid dynamics problems require the ability to model fluid transport on domains that do not conform to a mesh. Examples include bubble and droplet transport, mold filling, and general free surface flows that involve large interfacial motion and topological evolution. Conformal techniques like Arbitrary Lagrangian Eulerian (ALE) methods are subject to mesh tangling and preclude critical phenomena like coalescence and breakup.

*Correspondence to: David R. Noble, Sandia National Laboratories, MS-0382, P.O. Box 5800 Albuquerque, New Mexico 87185-0382, U.S.A.
†E-mail: drnoble@sandia.gov

Contract/grant sponsor: U.S. Department of Energy; contract/grant number: DE-AC04-94AL85000

Copyright © 2009 John Wiley & Sons, Ltd.

Recently, eXtended Finite Element Methods (XFEM) [1] and Generalized Finite Element Methods [2] have been developed for modeling physics on interfaces that do not coincide with the finite element mesh geometry. These methods address the interfacial physics and discontinuities without requiring the user to generate a boundary-fitted mesh. XFEM methods have been developed for fluid mechanics that include the interfacial physics of Stokes flow past particles [3, 4] and surface tension [5, 6]. In order to allow a discontinuous description across the interface, these methods require some level of enrichment of the elements that span the interface. Additional unknowns are assigned to one or more of the mesh entities (elements, nodes, sides, or edges) that are associated with these interfacial elements. Additional equations are formulated for these unknowns, and quadrature techniques [7, 8] are developed to account for the resulting discontinuous interpolation.

In the present work, an alternate approach is taken where a background, non-conformal mesh is decomposed into elements that conform to the boundaries of the fluid domain. Enrichment takes place by adding nodes that lie on the interfaces. Essentially, this is a generalization of the Finite Element Method using a Cartesian Grid with Added Nodes (FEMCGAN) [9] to unstructured finite elements. The focus of FEMCGAN is on Cartesian grids, and in that context it has been considered undesirable compared with Immersed Finite Element (IFE) methods. This is because the added nodes increase the size of the linear system of equations and significantly modify the structure of the matrix, rendering it unsolvable by Cartesian grid solvers [9].

Here, the conformal decomposition concept is applied to unstructured meshes of triangular and tetrahedral finite elements. The method is termed the Conformal Decomposition Finite Element Method (CDFEM). Like XFEM, CDFEM can be used to solve a multi-material problem starting with a mesh that does not conform to the geometry of the materials. This is done by decomposing the background mesh into elements that do conform to each of the materials. The starting point for this decomposition is to generate one or more level set fields that are used to describe the material geometry on the background mesh. These level set fields define how the non-conformal elements are decomposed into conformal sub-elements.

A primary goal of this paper is to quantify the accuracy of CDFEM. This is a concern because the level set surface cuts arbitrarily through the background mesh, and the quality of the resulting conformal elements may be quite poor. In [9], the accuracy of the approximating space in FEMCGAN was proved using the properties of the sub-elements formed by cutting the two-dimensional Cartesian grid. No similar analysis is included here. Instead, algorithms are designed that seek to maximize the quality of the resulting conformal elements. Theoretically, it is argued that the method is at least as accurate as XFEM methods in certain limits. The implementation of interfacial Dirichlet conditions is discussed since this has presented difficulties for XFEM approaches [10–12], but is handled trivially in CDFEM.

The accuracy of the method is quantified for multiple fluid problems using numerical examples. Specifically, simulations of potential flow in two and three dimensions are used to test the robustness and accuracy of the method in the context of Dirichlet and Neumann boundary conditions. Next, incompressible two-phase flow is considered to test the method for fluid–fluid interfaces. Last, flow about a periodic array of spheres is simulated, and the accuracy of the drag force is examined. These problems involve stationary fluid interfaces that do not move over the course of the simulation. This work is intended to verify the accuracy of the technique before applying it to moving interface problems in future work.

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

2. CONFORMAL DECOMPOSITION FINITE ELEMENT METHOD THEORY

2.1. Comparison to XFEM with Heaviside enrichment

Both CDFEM and XFEM can be used to describe problems with interfacial discontinuities. Figure 1 shows a comparison of the finite element descriptions that result when a piecewise linear triangle element is intersected by an interface (indicated by the red line) for both CDFEM and XFEM with Heaviside enrichment. In XFEM with Heaviside enrichment, a finite element field, u^h, can be expressed as

$$u^h(\mathbf{x}) = H(\phi(\mathbf{x}))\sum_i N_i(\mathbf{x})u_i^+ + (1 - H(\phi(\mathbf{x})))\sum_i N_i(\mathbf{x})u_i^- \tag{1}$$

where N_i are the standard nodal basis functions and $H(\phi)$ is the Heaviside function of the level set field, ϕ. Written this way, u_i^+ and u_i^- are the values of the nodal field for the positive and negative values of the level set field, respectively. Thus, with this form of enrichment, XFEM yields two independent piecewise linear descriptions such that the materials on each side of the interface can be independently described. The unknowns, however, remain at the vertices of the original non-conformal element. In this way, the material shown in blue and the material shown in green are described by separate piecewise linear approximations with basis functions derived from the standard basis functions of the non-conformal element. In comparison, CDFEM enriches the element by introducing new nodes where the interface intersects the edges of the non-conformal element. This produces a single piecewise linear approximation for the blue material, while the green material is described by two piecewise linear elements.

It is instructive to consider the conditions under which the CDFEM and XFEM descriptions are equivalent. In both methods, the blue material is described by a piecewise linear approximation. The two descriptions will be equivalent if the two interfacial nodes introduced in CDFEM are constrained to agree with the XFEM approximation. For the piecewise linear approximation considered here, that means that each interfacial node is constrained to be the same as the linearly

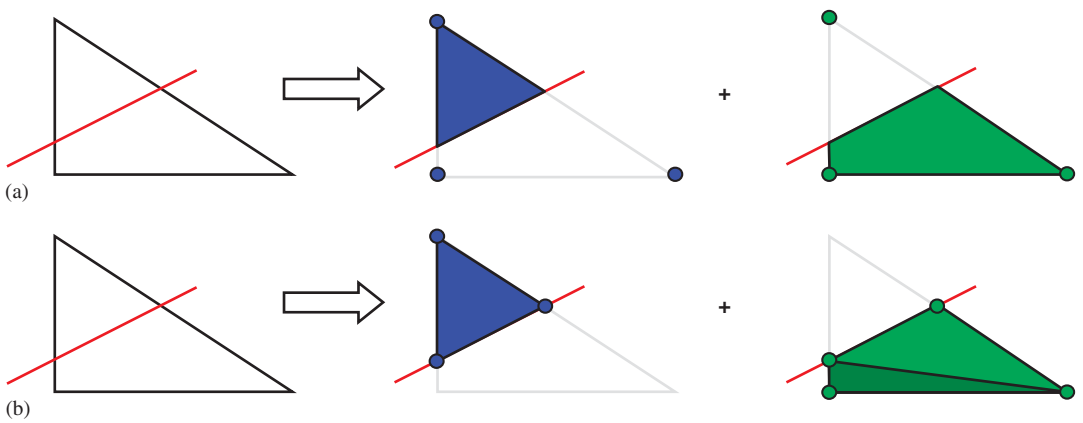

Figure 1. Comparison of finite element descriptions in CDFEM and XFEM with Heaviside enrichment: (a) independent piecewise linear approximations in XFEM with Heaviside enrichment and (b) piecewise linear elements produced in CDFEM.

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

interpolated edge value in XFEM. If this condition is satisfied, the CDFEM description and the XFEM description with Heaviside enrichment will be identical for the blue material. Likewise, the descriptions for the green material will be identical if the new interfacial nodes are constrained to lie in the space of the piecewise linear XFEM approximation for that material. While this discussion considered piecewise linear elements, the same conclusion may be drawn for any element type as long as the conformal elements are capable of replicating the same functions as the parent element.

In this discussion, it becomes clear that the discrete space introduced in CDFEM contains the space used in XFEM with Heaviside enrichment, since the XFEM space can be recovered by adding constraints on the nodes added in the conformal decomposition. According to [9], this property guarantees that the CDFEM accuracy is no less than that afforded by XFEM with Heaviside enrichment.

The relationship between these methods has interesting practical consequences. A primary motivation for CDFEM is to provide a layer of encapsulation separating all consideration of enrichment from the element assembly operations. This is not accomplished in XFEM where interpolation and quadrature are significantly modified to accommodate the enrichment. Based on the above arguments, an XFEM system can be assembled by first assembling the CDFEM description and then applying the constraints on the added nodes at the linear system level. This is very similar to the methods for handling constrained nodes in h-adaptivity. See [13] for a thorough discussion of multiple techniques for handling constrained nodes in the context of finite element assembly. Therefore, CDFEM provides a useful technique for assembling systems based on XFEM with Heaviside enrichment. By introducing the conformal mesh and the associated geometric constraints, the XFEM system can be formed using standard finite element assembly code with no modification whatsoever to account for the enriched description.

2.2. *Interfacial Dirichlet conditions*

The relationship between CDFEM and XFEM with Heaviside enrichment also has interesting ramifications for the application of interfacial Dirichlet conditions. This has proven to be a research issue in the context of XFEM methods [10–12]. In CDFEM, however, interfacial Dirichlet conditions are handled trivially by applying them directly to the interfacial nodes. So in one limit, CDFEM can recover the XFEM approximation by introducing appropriate constraints. On interfaces with Dirichlet conditions, however, these same nodes are subject to the Dirichlet constraints. The application of Dirichlet conditions in XFEM can therefore be considered to be the over-specification of the CDFEM problem. In this case, the XFEM constraints and Dirichlet constraints are being applied simultaneously, resulting in an over-constrained system. The consequences of this are the spurious oscillations encountered in XFEM that must be handled by selectively removing constraints, or by introducing stabilization or some other form of extra dissipation. None of these are required in CDFEM, where the Dirichlet constraints can be applied directly to the interfacial nodes. This is tested extensively in the numerical problems below.

3. CONFORMAL DECOMPOSITION FINITE ELEMENT METHOD IMPLEMENTATION

In CDFEM, one or more level set functions are used to describe the geometry on an unstructured, non-conformal mesh. These level set fields are then used to define how the non-conformal

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

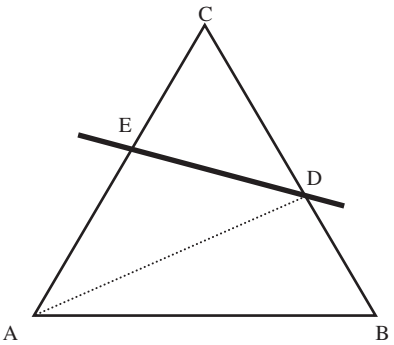

Figure 2. Angle-based decomposition of a non-conformal triangle into conformal triangles.

elements are decomposed into conformal sub-elements. The convention used here is to define interfaces where the level set changes sign. Up to 2^n materials can be described using n level set fields.

In this work, the level set field is assumed to consist of a piecewise linear field on triangular or tetrahedral elements. Thus, the interface consists of line segments in two-dimensions, and polygons in three-dimensions. Ignoring degenerate cases for a moment, the decomposition of a triangular element results in a triangular region and a quadrilateral region on opposing sides of the interface, as shown in Figure 2. To form a connected mesh of triangular elements requires decomposing the quadrilateral region into 2 or more conformal triangles. The choice of how this decomposition is performed will affect the quality of the resulting elements. In the sections below, strategies are developed for performing this conformal decomposition for both triangular and tetrahedral elements. First, however, the procedure for handling degeneracies is described.

3.1. Degeneracy handling

In general, the conformal decomposition must robustly handle degenerate cases. These occur whenever the interface passes exactly through a node. In addition, a robust scheme is needed for handling nearly degenerate cases. If nearly degenerate elements are not addressed, the resulting matrix system may be numerically singular. This issue was previously identified in context of XFEM with level sets [14].

The general procedure used to decompose the elements involves first cutting all edges that are crossed by the interface, followed by a triangulation step in which the sub-elements are formed from the resulting conformal edges. An edge is considered to be crossed if the level set values at the two nodes that support the edge have different signs. The sign of the level set is considered to be 'negative' if its value is less than zero. It is considered to be 'positive' if it is greater than *or* equal to zero. This choice for how to handle nodes that are exactly zero is clearly arbitrary, but they must be handled consistently. Every crossed edge is processed and a node is generated at the crossing. The coordinates of this node are calculated by linear interpolation on the edge. For an edge with nodal level set values of ϕ_1 and ϕ_2, the interface node is placed at

$$\mathbf{x}_i = \mathbf{x}_1 + \alpha(\mathbf{x}_2 - \mathbf{x}_1) \tag{2}$$

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

where $\mathbf{x}_1$ and $\mathbf{x}_2$ are the coordinates of the nodes and the position, α is found from linear interpolation:

$$\alpha = \frac{\|\phi_1\|}{\|\phi_1 - \phi_2\|} \tag{3}$$

If, however, $\alpha < \varepsilon$ or $\alpha > 1 - \varepsilon$, a new node is not generated. Instead, the interface is considered to pass through the nearest node of the edge. In this work, we take $\varepsilon = 10^{-4}$. This algorithm is essentially the same as that used in [14], where the level set value at the nearest node was set to zero in these degenerate cases. A detailed study of the sensitivity to this decomposition parameter has not been performed. It is expected that selecting too large of a value could cause significant errors due to the disagreement between the prescribed geometry and the decomposed geometry. On the other hand, too small of a value could cause multiple nodes to be created that are numerically coincident. Neither of these symptoms were encountered in any of the simulations for the chosen value.

In the subsequent triangulation step, the original vertex nodes and added edge nodes are connected into sub-elements. Because the edge nodes may be coincident with the vertex nodes, degenerate sub-elements may result, which can be identified by checking for duplicate nodes. Only non-degenerate sub-elements are retained in the decomposition.

3.2. Triangle decomposition

The description above shows that choices must be made in the process of obtaining a conformal decomposition, and that these decisions will affect the resulting element quality. In the two-dimensional case, at least two triangle elements will be formed from the quadrilateral region formed by cutting the element with a line. Additional triangles could be formed by adding one or more Steiner points, but that is avoided here to minimize the number of sub-elements formed by the decomposition. In graph theory and computational geometry, Steiner points are additional vertices that are not part of the input.

Therefore, the decomposition strategy for non-degenerate triangles comes down to the choice of which pair of opposing vertices to connect to break the quadrilateral region into triangles. For the case shown in Figure 2, this means choosing whether to connect AD or BE. The strategy taken here is to choose the diagonal that cuts the largest interior angle of the quadrilateral, shown by the dotted line in the figure. This is motivated by the work of Babuška and Aziz [15], which showed that the accuracy of finite element solutions degrades seriously as angles approach 180°, but the same is not true as angles approach 0°.

3.3. Tetrahedron decomposition

The decomposition of non-conformal tetrahedra into conformal sub-elements is more complicated than the two-dimensional case. When a plane cuts a tetrahedron, different subdomain topologies are generated depending on the number of nodes that lie on each side of the plane, as shown in Figure 3. An even more difficult issue is that the sub-element decomposition must take care to provide compatible discretizations across element faces.

All tetrahedra that are crossed by the interface plane either have three nodes of one sign and one node of the other sign, or have two nodes of each sign. In the first case, shown in Figure 3(a), the lone node, together with the three interfacial nodes formed by cutting the crossed edges, forms a conformal tetrahedron. The three other vertex nodes with the three interface nodes create a wedge shaped region that must be further decomposed into tetrahedra. For the second case, shown in

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

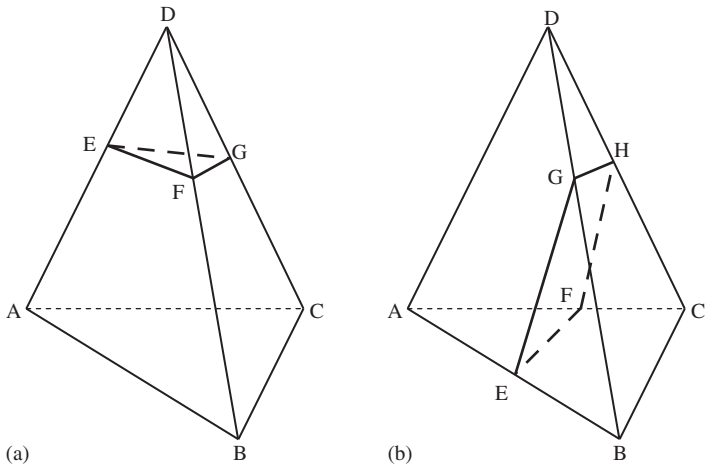

Figure 3. Isosurfaces in a linear tetrahedron: (a) three nodes of one sign (A, B, C) and one node with the other sign (D) and (b) two nodes of each sign (A, D and B, C).

Figure 3(b), each side of the interface consists of two vertex nodes and four interface nodes to form a wedge-shaped region. So the algorithm for decomposing a non-conformal tetrahedron into conformal tetrahedra consists of determining the signs of the vertex nodes, adding interface nodes to the crossed edges, and then decomposing any wedges into conformal tetrahedra.

It is this step of decomposing the wedges that must consider the requirement of forming compatible discretizations across element faces. A wedge decomposition is shown in Figure 4(a). Each of the three quadrilateral faces of the wedge must be decomposed by choosing which pair of opposing vertices to connect to generate two triangles. To provide compatible discretizations, the algorithm must decompose each face identically when it is processed from either of the owning elements. This requirement of compatible discretization across faces would naturally lead to a face-based criterion for selecting the diagonal. Ideally, we would like to use the same algorithm discussed for the two-dimensional decomposition. The problem is that not all combinations of choices are tetrahedralizable. Figure 4(b) shows a Schonhardt's polyhedron, which is not tetrahedralizable without the addition of one or more Steiner points [16]. Adding points is avoided because of the large number of sub-elements that would be produced. Any face-based criterion, including one based on the length of the diagonals, will suffer this same issue.

Instead, a partially node-based algorithm is used that has proved simple and robust across numerous tested applications. Node-based algorithms are guaranteed to provide a tetrahedralizable polyhedron, since the algorithm will result in diagonals on two of the faces of the wedge to emanate away from the same node. It is the lack of this feature that leads to the Schonhardt's polyhedron. The algorithm used in this work distinguishes between faces that lie on the interface and those that do not. For interfacial faces, the face-based largest angle criterion is used to select the diagonal. For the remaining faces, the decision is based on the magnitude of the level set function at each of the vertices. On these faces, the diagonal is selected such that they connect the node with the largest level set magnitude to the node opposite it. The rational for this algorithm is that it will preferentially choose to form new edges that are not aligned with the interface. While this algorithm has been shown to provide accurate and robust results, it is a place for possible

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

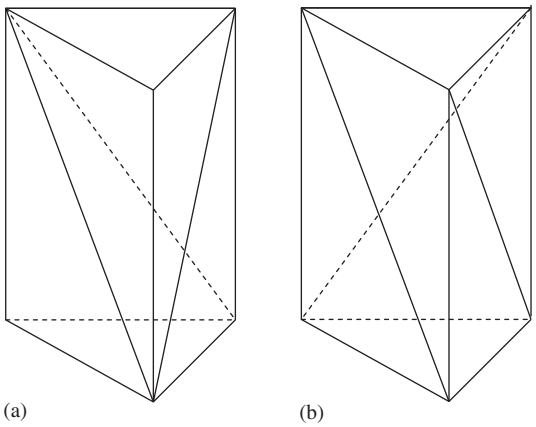

Figure 4. Decomposition of a wedge into tetrahedra: (a) a wedge that may be decomposed into 3 tetrahedra and (b) Schonhardt's polyhedron, which is not tetrahedralizable without the addition of Steiner points.

improvement in the future. If points are to be added to improve the quality of the decomposition, it is not clear whether these points should be added to the interior of the wedge or to the faces.

The above algorithm must be modified slightly to handle the case where multiple nodes have the same level set magnitude to within the precision of the machine. In these cases, the selection is based on some other nodal property. Ideally, an integer quantity like global node identifier would be used. If such a quantity is not available (as in parallel decompositions), we have successfully used the nodal coordinates as a selection criterion. Because the nodes have unique coordinates, this test is guaranteed to provide a unique identifier, resulting in a compatible discretization.

4. RESULTS

The proposed conformal finite element decomposition technique is tested on three steady-state problems from fluid dynamics: potential flow about a cylinder and sphere; steady viscous two-phase flow between concentric cylinders; and three-dimensional incompressible flow about a periodic array of spheres.

4.1. Two-dimensional potential flow about a cylinder

When posed in terms of the velocity potential, ψ, potential flow results in a Laplacian equation subject to Dirichlet and Neumann boundary conditions:

$$\nabla \cdot \nabla \psi = 0 \tag{4}$$

This provides a good test of the method since the application of Dirichlet conditions on embedded boundaries has proven to be a research issue [10–12]. The starting mesh is a non-conformal, unstructured mesh of triangular elements, shown in Figure 5(a). Two distance functions are generated, one for describing the cylinder surface, the other to describe an outer concentric circular ring that defines the extent of the computational domain. The radii of the inner and outer surfaces are 0.15 and 0.45, respectively, while the non-conformal mesh has unit extents in each direction. The

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

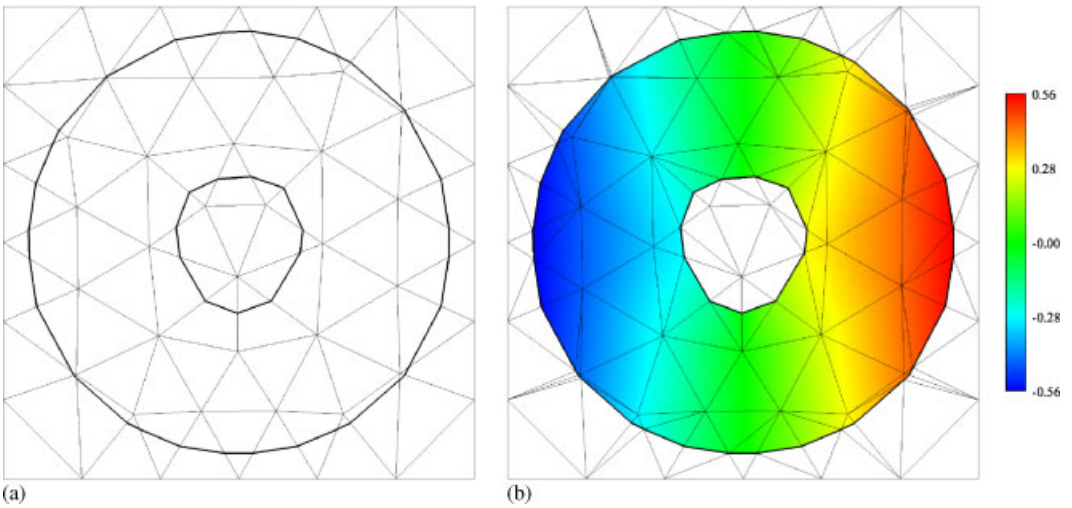

Figure 5. Decomposition of the non-conformal triangular mesh into a conformal one for the simulation of potential flow about a circular cylinder: (a) non-conformal mesh with level set surfaces indicated by heavy curve and (b) coarsest conformal mesh. Color indicates the magnitude of the velocity potential.

zero isosurfaces of the distance functions, which approximate the inner and outer radii, are also shown. These distance functions are calculated at each node using the exact distance to the nearest surface. The resulting isosurfaces only approximate the original circular boundaries because they are generated from the isosurfaces of the piecewise linear interpolant of the distance functions. The mesh that results from the conformal decomposition is shown in Figure 5(b). It is apparent that many of the conformal elements have small angles. A primary purpose of this study is to determine if optimal convergence rates can be obtained in spite of these low quality elements. Because the interface cuts the mesh arbitrarily, the mesh quality of the decomposed elements is not expected to improve with uniform refinement of the non-conformal mesh. Likewise, it is not expected to appreciably worsen with refinement. The relative volume of these low quality elements will decrease rapidly, however. For these reasons, the conformal mesh quality is not directly assessed in this study, but instead, the resulting accuracy is examined.

The surface of the cylinder is described by a zero normal flux, which is modeled using the natural finite element boundary condition. On the outer surface a Dirichlet condition is applied using the analytical solution evaluated at the surface nodes. The analytical solution for the velocity potential in polar coordinates (r, θ) is given by

$$\psi(r, \theta) = \left(1 + \frac{R_i^2}{r^2}\right) U r \cos(\theta) \tag{5}$$

where R_i is the cylinder radius. The magnitude of the velocity potential, U, is chosen to be a function of the cylinder radius and the domain radius, R_o, according to

$$U = \frac{R_o^2}{R_o^2 - R_i^2} \tag{6}$$

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

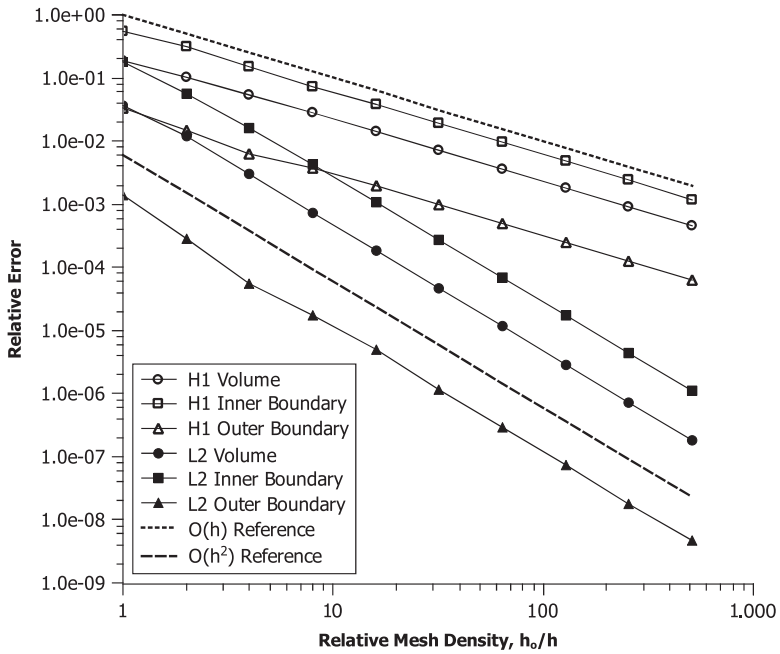

Figure 6. Convergence of potential flow about a circular cylinder.

The mesh shown in Figure 5(b) shows the coarsest mesh used in the convergence study. The color indicates the numerical solution for the velocity potential. The CDFEM solution for the potential flow is obtained using two modules within SIERRA mechanics, specifically KRINO and ARIA [17, 18]. First, the KRINO module initializes the level set field and performs the conformal decomposition of the initially non-conformal mesh. The ARIA module then solves a Laplacian equation for velocity potential subject to the applied boundary conditions.

Figure 6 shows the L_2 norms and H_1 seminorms of the error as a function of the mesh density, evaluated on the volume and on each of the cut surfaces. The coarse mesh is uniformly refined to provide progressively finer meshes. The relative mesh density is given by $h_o/h = 2^n$ where n is the number of levels of uniform refinement. All convergence rates are optimal for the linear triangle elements used in the simulation. The L_2 norms converge quadratically and the H_1 seminorms converge linearly. This is remarkable when compared with results reported in [12] for Dirichlet conditions on curved, embedded surfaces in XFEM.

Besides accuracy, another concern in CDFEM methods is the effect on the conditioning of the resulting system of equations. Because nearly degenerate elements are generated by the conformal decomposition, it is expected that ill-conditioning may result. This is tested by computing condition number estimates as the mesh is refined. The TRILINOS package [19] is used to solve the linear system while also providing an estimate of the condition number. The conjugate gradient method is used as a solver while a Lanczos matrix is formed to estimate the extreme eigenvalues. Figure 7 shows the estimated condition number for the CDFEM system with and without Jacobi preconditioning. For comparison, the estimated condition number for a simple conduction problem assembled on the uncut non-conformal mesh is also plotted. Both with and without Jacobi

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

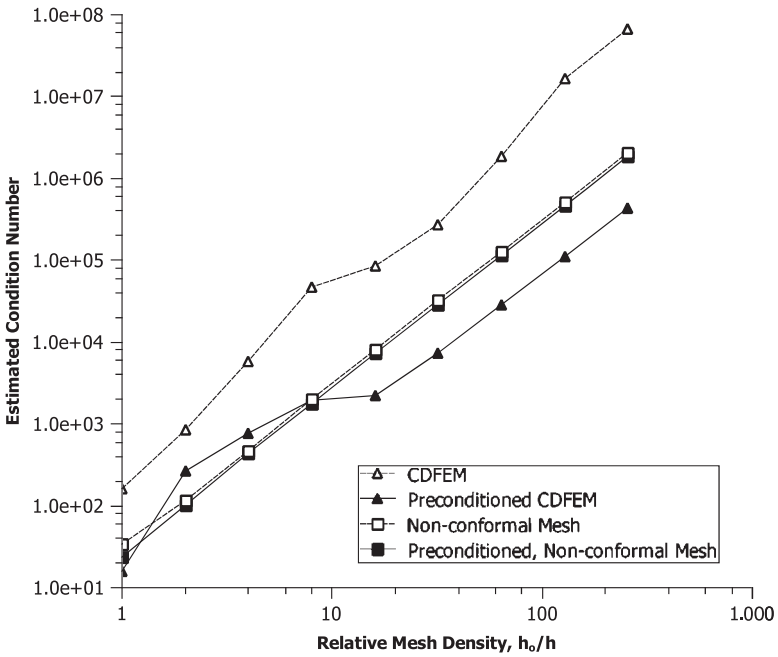

Figure 7. Estimated condition number for flow about a circular cylinder.

preconditioning, the simple conduction problem exhibits the expected $O(h^{-2})$ scaling for the condition number. While the CDFEM system without preconditioning exhibits a similar rate of growth, the condition number is more than 10 times higher. This confirms that conformal decomposition does indeed degrade the conditioning of the system of equations. This is true even though the CDFEM mesh has significantly fewer degrees of freedom, as the elements outside of the domain of interest have been removed. After applying simple Jacobi preconditioning, however, this ill-conditioning is essentially removed. On sufficiently fine meshes, the preconditioned system exhibits better conditioning than the conduction system on the uncut mesh. This result is consistent with the results from anisotropic mesh refinement [20] where it was found that simple diagonal preconditioning removed the ill-conditioning due to mesh degeneracy, leading to the same asymptotic growth in the condition number as arises for uniform mesh refinement. In practice, this means that the system of equations produced in CDFEM are readily solved using standard, iterative solvers.

4.2. Three-dimensional potential flow about a sphere

Similarly, the potential flow about a sphere is computed. Again Neumann conditions are used on the surface of the body while Dirichlet conditions are applied along the outer boundary, using the analytical solution. For this case, the analytical solution for the velocity potential, ψ, in polar coordinates (r, θ) is given by

$$\psi(r,\theta) = \left(1 + \frac{R_i^3}{2r^3}\right) U r \cos(\theta) \tag{7}$$

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

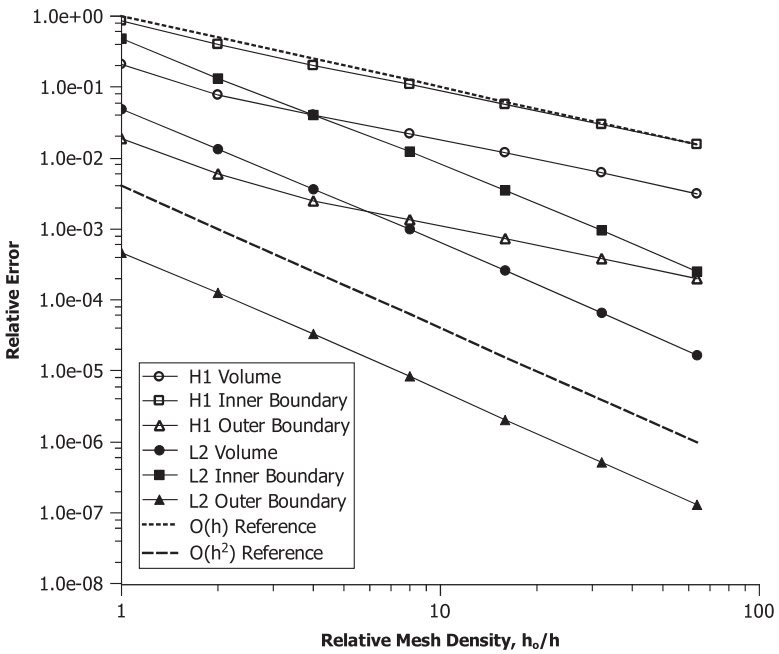

Figure 8. Convergence of potential flow about a sphere.

where R_i is the sphere radius. Here, the magnitude of the velocity potential, U, is chosen to be a function of the sphere radius and the domain radius, R_o, according to

$$U = \frac{R_o^3}{R_o^3 - R_i^3} \tag{8}$$

The radii are given by $R_i = 0.15$ and $R_o = 0.45$, and the non-conformal mesh has unit extents in each direction.

A plot of the resulting convergence is shown in Figure 8. Once again, optimal convergence rates are obtained. It is interesting that these convergence rates are obtained even though the resulting angles of the faces of the conformal elements were not optimized. Likewise, the dihedral angles were not considered in the decomposition. It is unknown if alternate strategies might be even more effective at lowering the approximation error.

4.3. Two-dimensional, two-phase, incompressible, viscous flow

The potential flow examples show the ability to apply different types of boundary conditions along embedded interfaces. This example shows the ability to capture the effect of material discontinuities across an interface in the context of incompressible fluids. The two-phase flow between co-rotating concentric cylinders is simulated. Figure 9 shows the mesh and example numerical solution with the color indicating the magnitude of the x-velocity. Again a non-conformal mesh has been decomposed, this time into 4 regions. The two annular regions are retained for the two fluids.

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

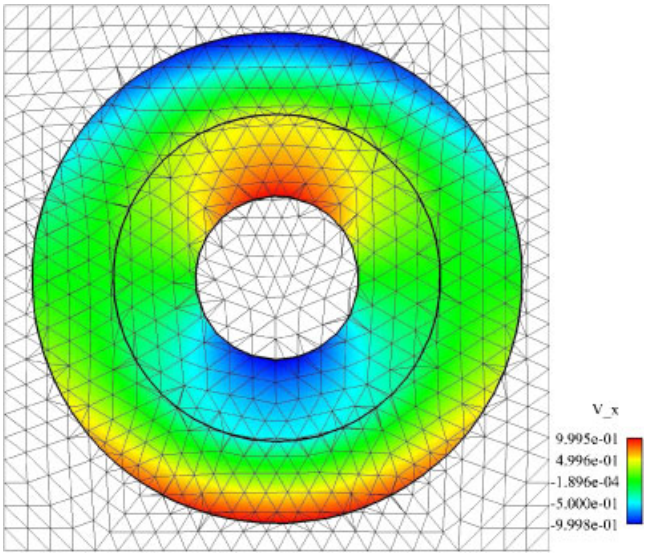

Figure 9. Decomposed mesh for one of the simulations of two-phase flow in the annular space between concentric cylinders. Color indicates the magnitude of the x-velocity.

The outer boundary is rotated counter-clockwise, while the inner boundary is rotated clockwise. The analytical solution for the azimuthal velocity for the inner fluid is given by,

$$v_\theta(r) = \frac{R_\Gamma^2\Omega_\Gamma - R_i^2\Omega_i}{R_\Gamma^2 - R_i^2} r + \frac{R_\Gamma^2 R_i^2(\Omega_\Gamma - \Omega_i)}{R_\Gamma^2 - R_i^2}\frac{1}{r} \quad \text{for } R_i \leqslant r \leqslant R_\Gamma \tag{9}$$

while for the outer fluid it is given by,

$$v_\theta(r) = \frac{R_o^2\Omega_o - R_\Gamma^2\Omega_\Gamma}{R_o^2 - R_\Gamma^2} r + \frac{R_o^2 R_\Gamma^2(\Omega_o - \Omega_\Gamma)}{R_o^2 - R_\Gamma^2}\frac{1}{r} \quad \text{for } R_\Gamma \leqslant r \leqslant R_o \tag{10}$$

where R_i is the radius of the inner cylinder, R_o is the radius of the outer cylinder, and R_Γ is the radius of the interface. Here,the radii are given by $R_i = 0.15$, $R_\Gamma = 0.3$, and $R_o = 0.45$, and the non-conformal mesh has unit extents in each direction. The angular velocity of the inner and outer cylinders are $\Omega_i = -1/R_i$ and $\Omega_o = 1/R_o$, respectively, so that the magnitude of the velocity on each cylinder is unity. The angular velocity of the interface, Ω_Γ, is a function of the inner and outer fluid viscosities, μ_i and μ_o, and is given by

$$\Omega_\Gamma = \frac{\mu_i R_i^2(R_o^2 - R_\Gamma^2)\Omega_i + \mu_o R_o^2(R_\Gamma^2 - R_i^2)\Omega_o}{\mu_i R_i^2(R_o^2 - R_\Gamma^2) + \mu_o R_o^2(R_\Gamma^2 - R_i^2)} \tag{11}$$

Dirichlet conditions are used on both the inner and outer surfaces to enforce this motion. The inner fluid is 4 times more viscous than the outer fluid, so a jump in the velocity gradient exists at the fluid–fluid interface. After the conformal decomposition, this gradient jump is contained in the basis of the resulting conformal sub-elements. This is not true for non-conformal elements.

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

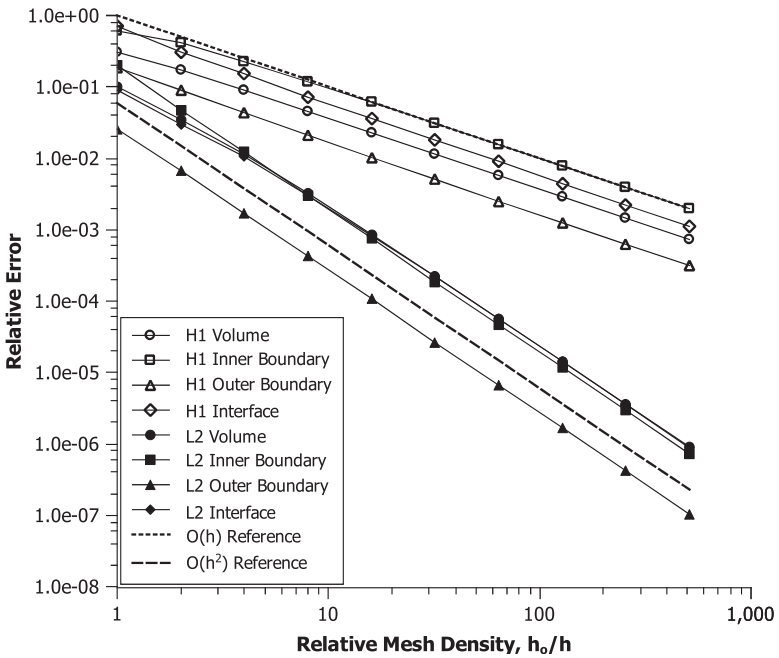

Figure 10. Convergence of incompressible, two-phase flow in an annulus.

The ARIA module within SIERRA mechanics solves the incompressible Stokes equations for the velocity field within the annulus. Since piecewise linear elements are used for both velocity and pressure, stabilization is required. Here, a relatively new stabilization method is employed that utilizes a projection method to circumvent the BB or inf–sup compatibility condition [21]. This Polynomial Pressure Projection Stabilization (PPPS) is significantly simpler to implement than the Pressure Stabilizing Petrov–Galerkin (PSPG) and contains no adjustable parameters.

Figure 10 shows the L_2 and H_1 seminorms of the error on the volume and on the inner, outer, and interface surfaces. The H_1 seminorm for the interface is evaluated from the elements within the inner annulus that touch the interface. Optimal convergence rates are obtained in each norm. The ability of CDFEM to capture the discontinuous velocity gradient is demonstrated in Figure 11, which shows the radial gradient of the azimuthal velocity component, $\partial v_\theta/\partial r$, as a function of radius. The curve indicates the analytical solution, which exhibits a jump at the fluid–fluid interface. The numerical points are obtained by plotting the piecewise constant value on each conformal element as a function of the radial position of the centroid of the element. The non-conformal mesh used in this simulation was generated by applying five levels of uniform refinement to the mesh shown in Figure 5(a). While the deviation from the analytical solution increases slightly near the inner cylinder and the fluid–fluid interface, the sharp jump in the gradient is clearly captured with no ringing or diffusive artifacts. This example demonstrates the ability of CDFEM to accurately capture the weak discontinuities associated with discontinuous material properties.

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

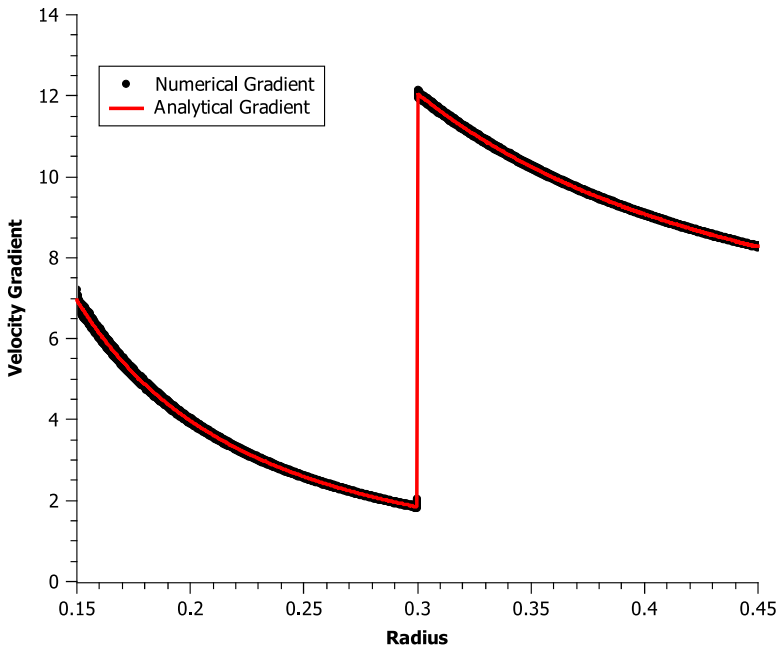

Figure 11. Discontinuous velocity gradient, $\partial v_\theta/\partial r$, in two-phase flow between concentric cylinders.

4.4. Three-dimensional, stokes flow about a periodic array of spheres

The final application is that of the Stokes flow about a simple cubic packing of spheres. The flow in a unit cell is computed with the original non-conformal domain decomposed into one that conforms to the sphere. The flow is symmetric across boundaries that are tangent to the flow direction. This is imposed with a Dirichlet condition zeroing the component normal to these boundaries. For the inlet and outlet, a pressure boundary condition is imposed weakly on the component of velocity in the flow direction by adding the contribution

$$\mathbf{q} = -\mathbf{p}_\circ \mathbf{n} + \mu \mathbf{n} \cdot \nabla \mathbf{v}^t \tag{12}$$

This boundary condition adds back the stress contribution for inlet/outlet flows where the pressure, $\mathbf{p}_\circ$, is prescribed. Here, μ is the viscosity, $\nabla \mathbf{v}$ is the velocity gradient, and $\mathbf{n}$ is the outward normal to the boundary. This boundary condition is used in open-flow applications to set the pressure datum. The assumption of a fully developed profile is implicit in this expression where $\mathbf{n} \cdot \nabla \mathbf{v} = 0$ has been used. The other two components of velocity are zeroed via Dirichlet conditions. No-slip is applied on the sphere for all three components of the velocity with Dirichlet conditions.

Solid fractions ranging from 0.001 to the simple cubic packing limit of 0.5236 are simulated. The Stokes equations are again solved using ARIA, with the same stabilization used for the previous test. The drag force is not evaluated using the pressure and velocity gradient due the low accuracy of the velocity gradients. Instead, the momentum equation residuals are evaluated prior to the applications of the Dirichlet boundary conditions. These residuals represent the nodal tractions. A much better approximation of the total drag force is obtained by summing these nodal forces. In Figure 12, the drag force, non-dimensionalized by the analytical value for a single sphere,

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

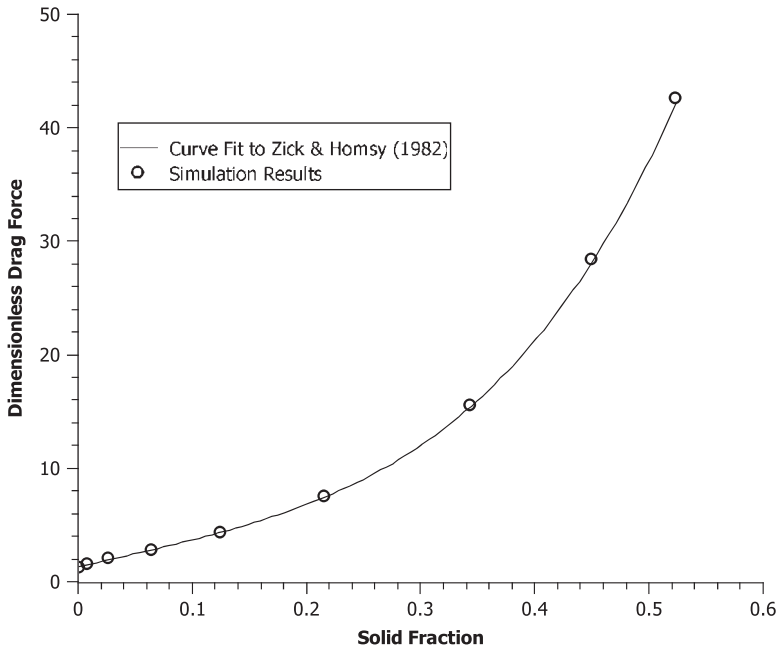

Figure 12. Drag force for flow about a periodic array of spheres as a function of solid fraction.

is shown compared against results obtained by Zick and Homsy [22] using a boundary integral technique. Very good agreement is obtained across the entire range of solid fractions. To quantify the agreement, a convergence study was done for a solid fraction of 0.125. For this packing, Zick and Homsy predict a dimensionless drag force of 4.292. Figure 13 shows the difference between the CDFEM results and this prediction as the mesh is refined. For comparison a reference line is shown for quadratic convergence. This comparison does not provide a formal quantification of the error since the Zick and Homsy prediction also includes numerical truncation error. Nonetheless, the convergence rate appears to be approaching second order for this drag calculation. Richardson extrapolation of the force predicted on the two finest grids yields a value of 4.292.

5. SUMMARY

In this work, a method is developed for simulating transport processes on domains defined by one or more level set functions on an unstructured mesh of triangular or tetrahedral elements. The method is termed the Conformal Decomposition Finite Element Method because the non-conformal mesh is decomposed into elements that conform to the level set description. Enrichment takes place by creating new nodes where the edges of the mesh are intersected by the level set interfaces.

CDFEM is closely related to XFEM [1] with Heaviside enrichment. In fact, the method is shown to be at least as accurate as XFEM. In addition, CDFEM readily accommodates Dirichlet boundary conditions without requiring stabilization or constraint reduction, as is needed in XFEM. Dirichlet conditions on curved surfaces remain robust and converge optimally over several decades of error

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

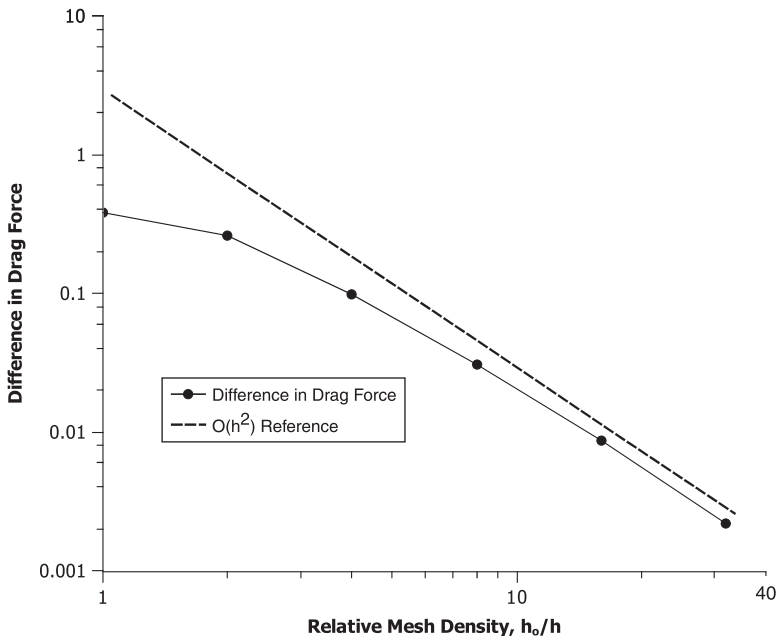

Figure 13. Convergence of drag force for flow about a periodic array of spheres with a solid fraction of 0.125.

reduction. For problems where an XFEM description is desirable, it can be recovered from the CDFEM description by formulating simple constraints for the created nodes. While the nearly degenerate conformal elements are shown to detrimentally impact the conditioning of the system of equations, this ill-conditioning is removed by simple, Jacobi preconditioning. Consequently, the system of equations produced in CDFEM are readily solved using standard iterative solvers.

Future work will examine CDFEM for moving interfaces. The goal is to get the best features of interface capturing methods, i.e. level set methods, and moving mesh methods, i.e. Arbitrary Lagrangian Eulerian (ALE) methods.

ACKNOWLEDGEMENTS

The authors would like to thank Professor John Dolbow for many useful discussions on the merits and applications of this work. This work was performed at Sandia National Laboratories. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

REFERENCES

1. Belytschko T, Moës N, Usui S, Parimi C. Arbitrary discontinuities in finite elements. *International Journal for Numerical Methods in Engineering* 2001; **50**:993–1013.
2. Strouboulis T, Copps K, Babuška I. The generalized finite element method. *Computer Methods in Applied Mechanics and Engineering* 2001; **190**:4081–4193.
3. Wagner GJ, Moës N, Liu WK, Belytschko T. The extended finite element method for stokes flow past rigid cylinders. *International Journal for Numerical Methods in Engineering* 2001; **51**:293–313.

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld

4. Wagner GJ, Ghosal S, Liu WK. Particulate flow simulations using lubrication theory solution enrichment. *International Journal for Numerical Methods in Engineering* 2003; **56**:1261–1289.
5. Chessa J, Belytschko T. An extended finite element method for two-phase fluids. *Transactions of the ASME. Journal of Applied Mechanics* 2003; **70**:143–157.
6. Chessa J, Belytschko T. An extended finite element method and level sets for axisymmetric two-phase flow with surface tension. *International Journal for Numerical Methods in Engineering* 2003; **58**:2041–2064.
7. Ventura G. On the elimination of quadrature subcells for discontinuous functions in the extended finite-element method. *International Journal for Numerical Methods in Engineering* 2006; **66**:761–795.
8. Holdych DJ, Noble DR, Secor RB. Quadrature rules for triangular and tetrahedral elements with generalized functions. *International Journal for Numerical Methods in Engineering* 2008; **73**:1310–1327.
9. Li Z, Lin T, Wu X. New Cartesian grid methods for interface problems using the finite element formulation. *Numerische Mathematik* 2003; **96**:61–98.
10. Ji H, Dolbow JE. On strategies for enforcing interfacial constraints and evaluating jump conditions with the extended finite element method. *International Journal for Numerical Methods in Engineering* 2004; **61**:2508–2535.
11. Moës N, Bechet E, Tourbier M. Imposing Dirichlet boundary conditions in the extended finite element method. *International Journal for Numerical Methods in Engineering* 2006; **67**:772–793.
12. Mourad HM, Dolbow J, Harari I. A bubble-stabilized finite element method for Dirichlet constraints on embedded interfaces. *International Journal for Numerical Methods in Engineering* 2007; **69**:772–793.
13. Carey GF. *Computational Grids*: *Generation*, *Adaptation*, *and Solution Strategies*. CRC Press: Boca Raton, 1997.
14. Moës N, Gravouil A, Belytschko T. Non-planar 3D crack growth by the extended fininite element and level sets—part I: Mechanical model. *International Journal for Numerical Methods in Engineering* 2002; **53**:2549–2568.
15. Babuška I, Aziz AK. On the angle conditions in the finite element method. *SIAM Journal on Numerical Analysis* 1976; **13**:214–226.
16. Dompierre J, Labbe P, Vallet M-G, Camarero R. How to subdivide pyramids, prisms, and hexahedra into tetrahedra. *Proceedings of the 8th International Meshing Roundtable*, South Lake Tahoe, CA, U.S.A., 1999; 195–204.
17. Notz PK, Subia SR, Hopkins MM, Moffat HK, Noble DR. *Aria 1.5*: *User Manual*, *SAND2007-2734*, Sandia National Laboratories, 2007.
18. Edwards HC. Sierra framework for massively parallel adaptive multiphysics applications. *SAND2004-6277C*, Sandia National Laboratories, 2004.
19. Heroux MA, Bartlett RA, Howle VE, Hoekstra RJ, Hu JJ, Kolda TG, Lehoucq RB, Long KR, Pawlowski RP, Phipps ET, Salinger AG, Thornquist HK, Tuminaro RS, Willenbring JM, Williams A, Stanley KS. An overview of the trilinos project. *ACM Transactions on Mathematical Software* 2005; **31**:397–423.
20. Graham IG, McLean W. Anisotropic mesh refinement: the conditioning of Galerkin boundary element matrices and simple preconditioners. *SIAM Journal on Numerical Analysis* 2006; **44**:1487–1513.
21. Dohrmann CR, Bochev PB. A stabilized finite element method for the Stokes problem based on polynomial pressure projections. *International Journal for Numerical Methods in Fluids* 2004; **46**:183–201.
22. Zick AA, Homsy GM. Stokes-flow through periodic arrays of spheres. *Journal of Fluid Mechanics* 1982; **115**:13–26.

Copyright © 2009 John Wiley & Sons, Ltd.

DOI: 10.1002/fld