

Material order-independent interface reconstruction using power diagrams[‡]

Samuel P. Schofield^{1,*}, Rao V. Garimella¹, Marianne M. Francois²
and Raphaël Loubère³

¹*Mathematical Modeling and Analysis (T-7), Los Alamos National Laboratory, Los Alamos, NM, U.S.A.*

²*Continuum Dynamics (CCS-2), Los Alamos National Laboratory, Los Alamos, NM, U.S.A.*

³*Mathematical Institute of Toulouse UMR-5219, Toulouse, France*

SUMMARY

We have developed a new, multi-material, piecewise linear interface reconstruction method that correctly locates the position of each material in the mesh cell and matches the required volume fractions with no material ordering required. This is different from other volume tracking interface reconstruction methods in which an improper material ordering may result in materials being incorrectly located within the cell. The new method utilizes a type of weighted Voronoi diagram, known as a power diagram, to reconstruct the interface from approximate material locations derived either from a particle model or quadrature formula. It works on structured and general polygonal grids, for an arbitrary number of materials and can be naturally extended to three dimensions. Published in 2007 by John Wiley & Sons, Ltd.

Received 12 January 2007; Revised 30 April 2007; Accepted 12 May 2007

KEY WORDS: volume of fluid; interface reconstruction; material ordering; power diagram

1. INTRODUCTION

The effective management and capture of interfaces is essential for accurate and reliable simulation of multi-material and multi-phase flows. Due to their strict conservation of materials, volume-of-fluid (VOF) methods using interface reconstruction are widely used and will be the focus of this article. VOF methods do not explicitly track the interface between materials, but rather advect volume fractions which prescribe the material composition of each cell of the mesh. When the

*Correspondence to: Samuel P. Schofield, T-7 MS B284, Los Alamos National Laboratory, Los Alamos, NM 87545, U.S.A.

†E-mail: sams@lanl.gov

‡This article is a U.S. Government work and is in the public domain in the U.S.A.

Contract/grant sponsor: US Department of Energy; contract/grant number: DE-AC52-06NA25396

Contract/grant sponsor: DOE Advanced Simulation and Computing (ASC) program

interface between materials is needed, the interface is recreated based on the material volume fraction in the cell and its surrounding cells [1–4].

A common problem impacting these reconstruction methods is their dependence on a specified material ordering, i.e. if more than two materials are present in a cell, the reconstruction may depend on the sequence in which the materials are processed. This is undesirable as it may improperly locate materials within the cell. In a finite volume implementation, this may result in material being incorrectly fluxed into neighbouring cells.

In this article, we propose and demonstrate a method that can reconstruct a multi-material interface with no dependence on material ordering. The method is very general: it works on unstructured grids, accommodates an arbitrary number of materials and extends naturally to three dimensions. The method utilizes a particle attraction model or approximate centroid calculation to infer the relative location of the materials in the cell. Using that information, the interface is reconstructed using a weighted Voronoi diagram, known as a power diagram, such that the required volume fractions are matched.

2. INTERFACE RECONSTRUCTION IN VOLUME-OF-FLUID METHODS

The VOF method, originally developed by Hirt and Nichols [5], advects the fractional volumes of each fluid in the cell to track materials in a flow simulation. The volume fraction, $f_m(\mathcal{C}_i)$, of a material, m , in a cell, \mathcal{C}_i of volume $\|\mathcal{C}_i\|$ is defined as

$$f_m(\mathcal{C}_i) = \frac{A_m(\mathcal{C}_i)}{\|\mathcal{C}_i\|} \quad (1)$$

where $A_m(\mathcal{C}_i)$ is the volume of the material in the cell. Early VOF methods used a simple interface that was defined to be a coordinate axis aligned line within each cell that partitioned the cell into the correct volume fractions. This is often referred to as the simple line interface calculation (SLIC) due to Noh and Woodward [6]. This interface structure was natural when combined with directionally split advection.

Youngs [7, 8] extended the method to permit the material interface to have an arbitrary orientation within the cell. In Youngs' method, the outward normal of the interface separating a material within a computational cell is taken to be the negative gradient of the volume fraction function, estimated using the volume fractions of that material in the neighbouring cells. On unstructured grids, the gradient of the volume fraction function is calculated with either a Green–Gauss formula [9] or a least-squares technique [3]. On either structured or unstructured grids, the interface is then defined by locating a line having the prescribed normal that cuts off the correct volume of material from the computational cell. This and other methods that allow an arbitrarily oriented linear interface within a cell are referred to as piecewise linear interface calculation (PLIC) methods. In general, the methods are first order and the interfaces will be discontinuous at cell boundaries. However, there are extensions that make the reconstruction second order using a local optimization [10] or interface smoothing [11, 12].

While PLIC methods work well for two-material cells, the method must be extended to reconstruct the interfaces between materials in computational cells with three or more materials present. In the 'onion-skin' approach, each material interface is assumed to separate two materials and consists of a single line segment with both endpoints on the computational cell boundary. This form of reconstruction works for simple layer structures only [7, 13]. A more general approach is




Figure 1. Nested dissection interface reconstruction for three materials: (a) the first material is removed leaving a smaller available polygon; (b) the second material is removed from the available polygon; (c) the remaining available polygon is assigned to material 3; (d) the resulting partitioning of the computational cell; and (e)–(g) show the same procedure but the materials are processed in a different order leading to a different reconstruction (h).

‘nested dissection’ [14], where each material is separated from the others in a specified order. This process is illustrated in Figure 1. In Figure 1(a) the first material is separated from the remaining materials leaving a pure material region and an available polygon. Then in Figure 1(b), the second material is removed from the available polygon according to its calculated normal and required volume. Finally, in Figure 1(c), the remaining available polygon is assigned to the last material. This same process can be used for an arbitrary number of materials. Unfortunately, the result of this method depends on the order in which the materials are processed. In Figures 1(e)–(g), the reconstruction of the cell with the same volume fractions and interface normals is performed, but with the materials processed in a different order, leading to a different reconstruction as is seen by comparing Figures 1(d) and (h).

The effects of material order dependency are further illustrated in Figure 2. With the correct material ordering used in Figure 2(b), the interface reconstructed by a first-order PLIC method, is close to the correct configuration shown in Figure 2(a). In this example, the reconstruction would be identical to the correct configuration using a second-order reconstruction method. Using an incorrect ordering results in substantial degradation of the interface as shown in the reconstructions in Figures 2(c) and (d).

The calculation of the interface normal is also affected by the presence of multiple materials. In Youngs’ method, the gradient of the volume fraction function gives the normal of the interface. In two-material simulations, the orientation of the interface is independent of which material is used to calculate the normal since the volume fractions satisfy the relationship $f_1 = 1 - f_2$ which gives $\nabla f_1 = -\nabla f_2$. However, for multi-material simulations this simple relationship between the volume fractions does not hold. Since the gradient or normal direction is calculated using the

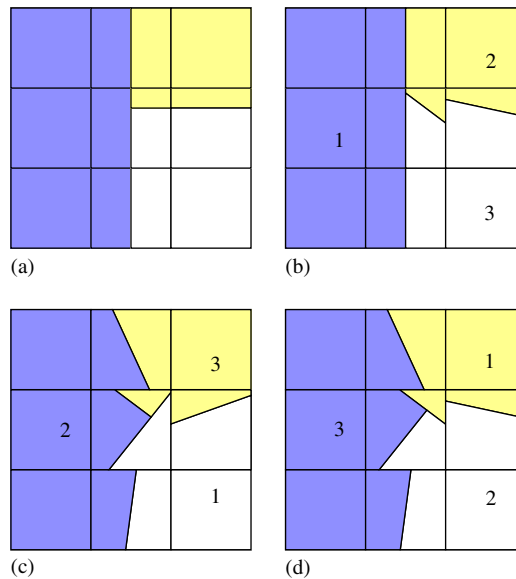


Figure 2. (a) Correct reconstruction and (b)–(d) first-order PLIC-based reconstruction using different material orderings. The numbers designate the order in which the materials were processed.

volume fraction data in a neighbourhood of the cell being reconstructed, it may happen that in a two-material cell, \mathcal{C}_i , $\nabla f_1(\mathcal{C}_i) \neq -\nabla f_2(\mathcal{C}_i)$ due to the presence of another material in the neighbourhood used for the gradient calculation. Two different reconstructions are possible depending on whether $\nabla f_1(\mathcal{C}_i)$ or $\nabla f_2(\mathcal{C}_i)$ is chosen as the direction for the interface normal. This introduces a material order dependence in even two-material cells. This is shown in Figures 2(c) and (d) as the interface of the two-material cell on the right side of the mesh has a different normal depending on which material was processed first.

These incorrect reconstructions may adversely impact the material advection in the simulation. If the advection routine is based on fluxing volumes through the edges of the cell, the improper material ordering may result in a material being fluxed early or not at all. For example, if the flow is moving diagonally towards the top right in Figure 2(c), the white material will move into the top right cell prematurely. This can lead to a breakup of the interface, a phenomenon known as ‘flotsam and jetsam’.

Selecting a global ordering can be problematic as the appropriate ordering for one region of the mesh may be quite wrong for another. To remedy this, there has been some work on deriving the material order. The geometrically derived material priority by Mosso and Clancy [15] is based on the assumption of a layer structure and works by approximating the local centre of mass of each material; then, based on the relative locations along a line, it selects an ordering. A similar approach was developed by Benson [16]. However, both methods can fail in the presence of a triple point.

In order to eliminate the material order dependence in multi-material interface reconstruction, we have developed a novel method consisting of two steps:

1. Relative locations of materials in multi-material cells are inferred using a particle model or quadrature formula.

2. Using these locations, the interface is reconstructed using a form of weighted Voronoi diagram known as a power diagram.

The method is completely general, working on general polygonal grids with an arbitrary number of materials in each cell. In addition, it can be naturally extended to three dimensions. In contrast to existing SLIC or PLIC methods, all materials are processed simultaneously and, as such, have no material order dependency. Furthermore, unlike methods such as the triple point method of Choi and Bussmann [13], no additional extensions are required to accommodate an arbitrary number of materials within the cell. The reconstruction will automatically give either the appropriate layer structure or multiple triple point configuration.

3. DETERMINATION OF MATERIAL LOCATIONS

We describe two methods for determining the approximate location of each material in the cell:

1. A particle model, where particles representing the materials evolve according to an attraction/repulsion model.
2. A method of approximately calculating the centre of mass of a material in a neighbourhood of the cell.

3.1. Material location with a particle attraction–repulsion model

In the first step of the method, a number of particles representing the materials are placed in multi-material cells and any pure or mixed neighbouring cells. A particle, P_i , has a position, \mathbf{x}_i , velocity $\mathbf{v}_i = d\mathbf{x}_i/dt$ and material $m(i)$, and is constrained to stay within the cell in which it is initially placed.

Taking inspiration from molecular dynamics [17, 18] and smoothed particle hydrodynamics [19, 20], we evolve the particle positions according to ‘forces’ based on the particles’ relative locations and materials. The positions of the particles are updated through time integration of a set of ordinary differential equations,

$$\frac{d\mathbf{x}_i}{dt} = \mathbf{V}_i \quad (2)$$

$$\mathbf{V}_i = \sum_{j:m(j)=m(i)} \mathbf{V}_{\text{att}}(\mathbf{x}_i, \mathbf{x}_j) + \sum_{j:m(j) \neq m(i)} \mathbf{V}_{\text{rep}}(\mathbf{x}_i, \mathbf{x}_j)$$

where \mathbf{V}_{att} and \mathbf{V}_{rep} are the prescribed attractive and repulsive ‘forces’ in the direction $\mathbf{x}_j - \mathbf{x}_i$. Particles of the same material attract each other until they are very close, at which point they start to repel each other. Particles of different materials repel each other. In our tests, the particles start at random locations within their cell, but they can be initialized using other means such as their relative locations in a cell at a previous time step.

The particle–particle ‘forces’ (plotted in Figure 3) are prescribed as

$$\mathbf{V}_{\text{att}}(\mathbf{x}_i, \mathbf{x}_j) = \begin{cases} -1, & d_{ij} < \delta \\ 1 - 2d_{ij}^4 + d_{ij}^8, & \delta \leq d_{ij} \leq 1 \\ 0, & d_{ij} > 1 \end{cases} \quad (3)$$




Figure 3. Particle attraction and repulsion 'forces' used in the model.

$$\mathbf{V}_{\text{rep}}(\mathbf{x}_i, \mathbf{x}_j) = \begin{cases} -(1 - 2d_{ij}^4 + d_{ij}^8), & d_{ij} \leq 1 \\ 0, & d_{ij} > 1 \end{cases} \quad (4)$$

where $d_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|/2.5h$ is the distance between points scaled by an interaction distance, taken to be 2.5 times the characteristic mesh size h , and $\delta = 0.05$. Unlike a traditional mechanical model, the 'forces' here actually prescribe the instantaneous velocities of the particles.

In a cell, \mathcal{C}_i , the number of particles, $N(\mathcal{C}_i)$ is

$$N(\mathcal{C}_i) = \left\lfloor N_p \times \frac{\|\mathcal{C}_i\|}{A_0} \right\rfloor \quad (5)$$

where N_p is a prescribed constant (usually around 30), $\|\mathcal{C}_i\|$ is the area of the cell, A_0 is a reference cell area for the grid (for example, on a uniform Cartesian grid, $A_0 = h^2$ where h is the grid spacing) and $\lfloor a \rfloor$ is the floor function giving the greatest integer less than or equal to a . Each particle has a designated material type, corresponding to a material present in the cell. Each material that is present in the cell is represented by the same number of particles, $N(\mathcal{C}_i)/N_m^i$, where N_m^i is the number of materials present in the cell. We found that making the number of particles representing each material proportional to the volume fraction of the material often leads to unsatisfactory results. If the volume fraction is small, the material will be represented only by a few particles, which are not sufficient to provide a reliable estimate of the location of the material within the cell. In addition, we found that for unstructured, general polygonal grids, making the number of particles proportional to the area of the cell was important. Otherwise, the particles tend to cluster in regions of the mesh with a concentration of smaller cells.

Once the particles are distributed, the particle model is run. Since the model prescribes instantaneous velocities and not true forces, the particles may remain in perpetual motion unless the system is forced to 'cool'. The velocity of each particle is rescaled at each time step to decrease the kinetic energy of the system and force the particles to settle into a final configuration. At time step n in the time integration of Equation (2), the kinetic energy of all the particles is

$$KE^n = \sum_i \frac{1}{2} \|\mathbf{V}_i\|^2 \quad (6)$$

After the system is sufficiently agitated, typically after 5 to 10 time steps, we force the kinetic energy to decrease as,

$$KE^{n+1} \leq \alpha KE^n \quad (7)$$

where $0 < \alpha < 1$. In practice, α is set to be 0.7–0.9. If $KE^{n+1} \geq KE^n$, all the particle velocities are scaled as

$$\mathbf{V}'_i = \sqrt{\alpha \frac{KE^n}{KE^{n+1}}} \mathbf{V}_i \quad (8)$$

To speed up the calculation, we use a variable time step with a new Δt calculated after each time step as

$$\Delta t = \frac{0.1}{2 \|\mathbf{V}_{\max}\|} \quad (9)$$

where $\|\mathbf{V}_{\max}\| = \max_i \|\mathbf{V}_i\|$ where \mathbf{V}_i is as defined in Equation (2).

The positions are then updated as

$$\mathbf{x}_i^{n+1} = \mathbf{x}_i^n + \Delta t \mathbf{V}'_i \quad (10)$$

If a particle goes outside the cell, it is placed back in the cell by repositioning it to the centre of the triangle formed by the old position, the new position, and the centre of the cell. If that fails, the particle is kept in its old position.

The particles are allowed to evolve for a number of time steps until the average kinetic energy of each particle has dropped below a specified stopping criterion. The particle model exhibits rapid convergence to the particle clusters, usually requiring under 20 time steps to converge to approximately the final positions. We have conducted statistical tests that show the model displays little sensitivity to the random initial particle positions, with standard deviations in the final material locations typically less than 5% of the mesh spacing [21].

Once the velocity of the particle has dropped to a sufficiently low threshold, the positions of the materials in the cell need to be derived from the final location of the particles. However, the particles of a material may form multiple groups. A clustering algorithm is needed to detect the multiple clusters and utilize that information to capture the subcell structure. A naive averaging of the particle positions for each material can yield reasonable results if each material in the cell is accurately described by a single convex polygon, but it will not detect the presence of multiple clusters of particles which may occur when thin filament-like structures are present.

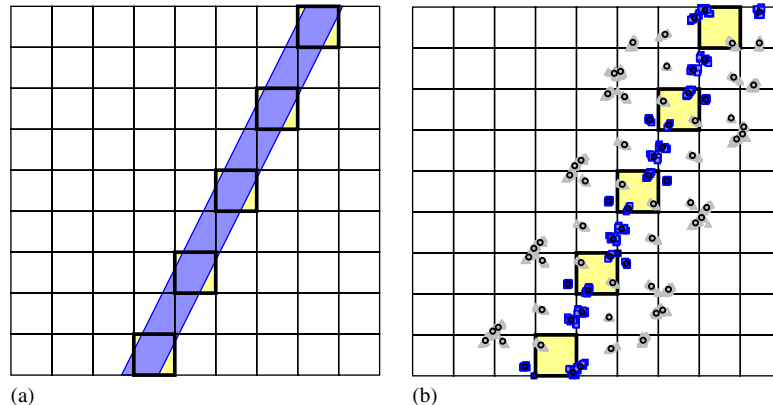


Figure 4. (a) Thin filament structure and (b) converged particle locations from the volume fraction data. The cluster locations are designated with the dark circles. The highlighted cells show where the filament has multiple interfaces within the cell. In these cells, away from the boundary, multiple clusters of particles are detected by the k -means clustering algorithm.

A simplified version of the k -means clustering algorithm [22], is used to detect if the particles of a material form one or two clusters. The algorithm proceeds as follows, for each material:

1. Initialize cluster 1 with the position of any particle of that material.
2. Initialize cluster 2 with the position of the particle of the same material farthest from cluster 1.
3. Assign each particle to the closest cluster.
4. Compute the mean position of all particles in each cluster to obtain new cluster locations.
5. Repeat steps 3 and 4 until the cluster locations converge or one cluster contains no particles.
6. If the two clusters are sufficiently close, that is, the distance between them is less than 10% of the characteristic mesh size, then use the average location of all of the particles of that material instead.

The clustering algorithm converges rapidly and is robust.

Figure 4 shows the example of a narrow filament. The particles reflect the thin filament structure present. The clustering algorithm is able to detect the presence of the multiple clusters within a cell. With the exception of cells at the boundary, multiple clusters are present and detected in all of the cells where the filament has two interfaces within the cell. These cells are indicated as the highlighted cells in Figure 4.

If there are multiple clusters of a material, the volume of material must be partitioned between the clusters. The number of particles in each cluster does not give a reliable means to do this as it can vary depending on the initial conditions. Instead, we partition the material volume equally between the clusters as we have no knowledge of the relative volume of the two regions of the material within the mesh cell.

3.2. Material location via approximate centres of mass

An alternative to the particle model is the direct calculation of an approximate centre of mass of each material in a subset of the mesh around the cell being reconstructed. In a region of the mesh,

Ω , assume a material occupies a non-trivial subset $\Omega_m \subseteq \Omega$. The centre of mass of the material in the mesh region, Ω is given as

$$\mathbf{x}_m(\Omega) = \frac{\int_{\Omega} \mathbf{x} \chi_m(\mathbf{x}) \, d\mathbf{x}}{\int_{\Omega} \chi_m(\mathbf{x}) \, d\mathbf{x}} \tag{11}$$

where $\chi_m(\mathbf{x})$ is the characteristic function for material m , defined as

$$\chi_m(\mathbf{x}) = \begin{cases} 1, & \mathbf{x} \in \Omega_m \\ 0, & \mathbf{x} \notin \Omega_m \end{cases} \tag{12}$$

However, Ω_m and χ_m are unknown.

If we assume that within each cell the centre of mass of each material contained within that cell is approximately the centroid of the mesh cell, then Equation (11) may be approximated for a region of the mesh larger than the cell of interest. In particular, if we take Ω to be the mesh cell we wish to reconstruct and all of its neighbours, the approximation takes the form

$$\tilde{\mathbf{x}}_m(\mathcal{C}_i) = \frac{\sum_{\mathcal{C}_j \in \mathcal{N}_i} f_m(\mathcal{C}_j) \|\mathcal{C}_j\| \mathbf{x}_c(\mathcal{C}_j)}{\sum_{\mathcal{C}_j \in \mathcal{N}_i} f_m(\mathcal{C}_j) \|\mathcal{C}_j\|} \tag{13}$$

where \mathcal{N}_i is the set of mesh cells including the cell being reconstructed, \mathcal{C}_i , and all of its neighbouring cells. $\mathbf{x}_c(\mathcal{C}_j)$ is the centroid of mesh cell \mathcal{C}_j . Similar formulations for a regional centre of mass derived from the volume fraction data have been used by a number of investigators [15, 16, 23].

These approximate centres of mass will be for the region \mathcal{N}_i and will in general not be contained within the cell \mathcal{C}_i . However, it gives some information regarding the relative location of materials in a cell.

4. POWER DIAGRAM-BASED INTERFACE RECONSTRUCTION

Once the materials are located in a cell, the interfaces within the cell, separating the materials, are constructed using a power diagram. A power diagram or Laguerre diagram [24, 25] is a generalization of a Voronoi diagram generated from a set of points, S , each with an associated radius or weight. In this context, the generators will be the points determined either by the particle clusters or the approximate centroids. The Laguerre distance from a point $\mathbf{x} \in \mathbb{R}^n$ to a point mass, $s_i \in S$ with $s_i = (\mathbf{x}_i, w_i)$ is defined as

$$d_L^2(\mathbf{x}, s_i) = d^2(\mathbf{x}, \mathbf{x}_i) - w_i \tag{14}$$

where $d^2(\mathbf{x}, \mathbf{x}_i) = \sum_{i=1}^n (x - x_i)^2$ is the usual Euclidean distance in \mathbb{R}^n . If w_i is replaced with w_i^2 in Equation (14), the resulting distance is called the *power* of the point \mathbf{x} with respect to \mathbf{x}_i .

Each cell in the power diagram is the set of points

$$\text{cell}(s_i) = \{\mathbf{x} \in \mathbb{R}^n \mid d_L^2(\mathbf{x}, s_i) < d_L^2(\mathbf{x}, s_j) \ \forall s_j \in S, s_j \neq s_i\} \tag{15}$$

As with Voronoi diagrams, each power diagram cell is convex.

The weight associated with a point generator can be interpreted as the square of the radius of a circle centred at that point. The power bisector (a chordale in Aurenhammer’s terminology




Figure 5. Power diagram with four generators ($V1$ – $V4$) and their weight circles. The generators $V1$ and $V2$ do not lie within their corresponding cells.

[24]) between two points $s_i = (\mathbf{x}_i, w_i)$ and $s_j = (\mathbf{x}_j, w_j)$ is the line perpendicular to the segment connecting the points \mathbf{x}_i and \mathbf{x}_j and is located by finding a point, \mathbf{x}_0 on that segment such that $d_L^2(\mathbf{x}_0, s_i) = d_L^2(\mathbf{x}_0, s_j)$.

An example of power diagram is shown in Figure 5. Clearly, if all point masses have equal weight (or radius), the power diagram reduces to the usual Voronoi diagram. Unlike a Voronoi diagram, a point in the point set generating the diagram does not necessarily lie in the cell to which it corresponds, as is this case for generators $V1$ and $V2$ in Figure 5. Furthermore, the cell corresponding to a given point mass, may be trivial, that is $\text{cell}(s_i) = \emptyset$. In practice, this is not a problem. Imai *et al.* [25] provide a useful lemma that provides a sufficient condition for the power diagram cell of a point to be non-trivial:

Given the power diagram for a finite set of point masses, $S = \{s_1, \dots, s_n\}$, $\text{cell}(s_i)$ is non-trivial if s_i lies on a corner of the convex hull of S .

If only three materials are present (i.e. S consists of only three point masses), then the point mass corresponding to each material must necessarily be a corner of the convex hull assuming the three points are not collinear. This ensures that the power diagram in each cell with three materials will always be non-trivial for all choices of weights. If the points are collinear, then the cell is partitioned by two parallel lines which can obviously be made to cut off the appropriate volume fractions. The same argument applies to three and four material cells in 3D. Proving this for larger numbers of materials is more difficult. However, we have not encountered a scenario where we are unable to find a power diagram that matches the required volume fractions.

The power diagram can be constructed in a number of ways. A power diagram may be created through a randomized, incremental algorithm [26], similar to the incremental construction of a Delaunay triangulation. A Voronoi diagram of the point generators may also be efficiently converted into a power diagram [27]. However, we have chosen a simple algorithm that intersects all of the

mutual power bisectors. It has $\mathcal{O}(n^2)$ asymptotic complexity in the number of generator points. For our applications, the number of materials in a cell, corresponding to the maximum possible number of power diagram cells, is small (typically 5 or less), so the asymptotic complexity of the construction algorithm is not a problem. In addition, it has proven to be robust in finite precision arithmetic.

4.1. Matching volume fractions

The volume fractions of the materials in a cell are matched by iteratively adjusting the weights of the point generators, thereby adjusting the area of the power diagram cells once clipped to the mesh cell. This requires the solution of a set of non-linear equations

$$A_m(\omega_1, \dots, \omega_{N_m}) = \|\mathcal{C}_i\| f_m, \quad m = 1, \dots, N_m \quad (16)$$

where $A_m(\omega_1, \dots, \omega_{N_m})$ is the area of the power diagram corresponding to material m after it has been clipped by the bounding mesh cell polygon with area $\|\mathcal{C}_i\|$. $f_m(\mathcal{C}_i)$ is again the volume fraction for material m in cell \mathcal{C}_i . The constraint

$$\sum_m^{N_m} A_m(\omega_1, \dots, \omega_n) = \|\mathcal{C}_i\| \quad (17)$$

reduces the number of equations to $N_m - 1$. Specifically, this is done by forcing one of the weights to be a specified value. This enforces a unique set of weights for the desired power diagram.

A Newton procedure with a finite-difference Jacobian is used to solve Equations (16) and (17). Some caution is required, since the area of each cell is bounded above and below, that is

$$0 \leq A_m(\omega_1, \dots, \omega_n) \leq \|\mathcal{C}_i\|, \quad m = 1, \dots, N_m \quad (18)$$

For extreme values of the weights, some of the power diagram cells will be outside of the mesh cell and as such have zero area once clipped to the mesh cell. Furthermore, the A_m will be flat (that is they have a zero gradient) making the Newton procedure fail. As a result, the Newton procedure needs to adjust for overshoots to make sure it does not end up in this region. This is simply done by reducing the size of the Newton step at each iteration if it exceeds those bounds. We found the procedure to be robust and efficient, typically matching the required volumes to within $10^{-12} \|\mathcal{C}_i\|$ in 3–6 iterations.

For the initial guess, we use equal weights for all the point generators if all of the generators lie within the mesh cell being reconstructed. If any of the generators are outside the cell, as may happen with the approximate centre of mass calculation, the initial weights are assigned such that the power bisectors between all the generators go through the centroid of the cell. This ensures that the initial power diagram will not have any cells outside of the mesh cell.

4.2. Reconstruction fidelity

The use of the power diagram to reconstruct the interface is based on the assumption that we have been able to obtain an approximation to the location of each material in the cell. Indeed, if the point generators for the power diagram are the exact centres of mass of the material subcells, the reconstruction is representative of the actual configuration. In Figures 6 and 7, the power diagram reconstruction of different three-material cells, using the actual material centres of mass

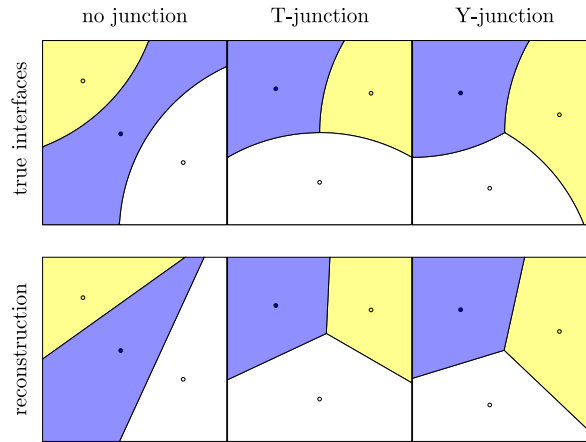


Figure 6. The true material interfaces and their centres of mass are shown in the top row with the power diagram reconstruction using the centres of mass in the bottom row. Here, the radius of curvature of all of the true interfaces is set to $R/h = 1$.




Figure 7. The asymptotic convergence of the power diagram reconstruction using the centres of mass as the radius of curvature of the interfaces becomes large. The power diagrams are unable to recreate the interfaces exactly. Here $R/h = 256$.

as generators, are shown. The power diagram provides a reasonable approximation to the true interface for interfaces with a radius of curvature comparable to the mesh spacing. As the radius of curvature goes to infinity, the interfaces straighten, but the power diagram does not necessarily converge to the true interface as illustrated in Figure 7. Still, the materials retain their relative locations within the cell.

In the following section, we present some static interface reconstruction examples using the power diagram-based reconstruction with point generators derived from the particle model and the approximate centroid calculation.

5. NUMERICAL EXPERIMENTS

To assess the overall performance of the method, a number of numerical experiments were performed.

In Figure 8, a four-material interface reconstruction on a structured grid is shown. As can be seen, our reconstructions respect the symmetry in the problem. The particle model derived material locations shown in Figure 8(a) are completely symmetric and result in a near perfect reconstruction. The reconstruction using approximate centre of mass-derived material locations shown in Figure 8(b) is exact. Observe that the material locations shown in the inset are outside of the centre cell. In the material order-dependent Youngs' method reconstructions in Figures 8(c) and (d), the reconstruction in the centre cell is not symmetric due to the material ordering. In fact, for this example, no material ordering will enable Youngs' method to create a completely symmetric reconstruction.

For the unstructured grid shown in Figure 9, neither power diagram-based reconstruction exactly reproduces a straight line, indicating that the reconstruction cannot be second order [10]. Still, in

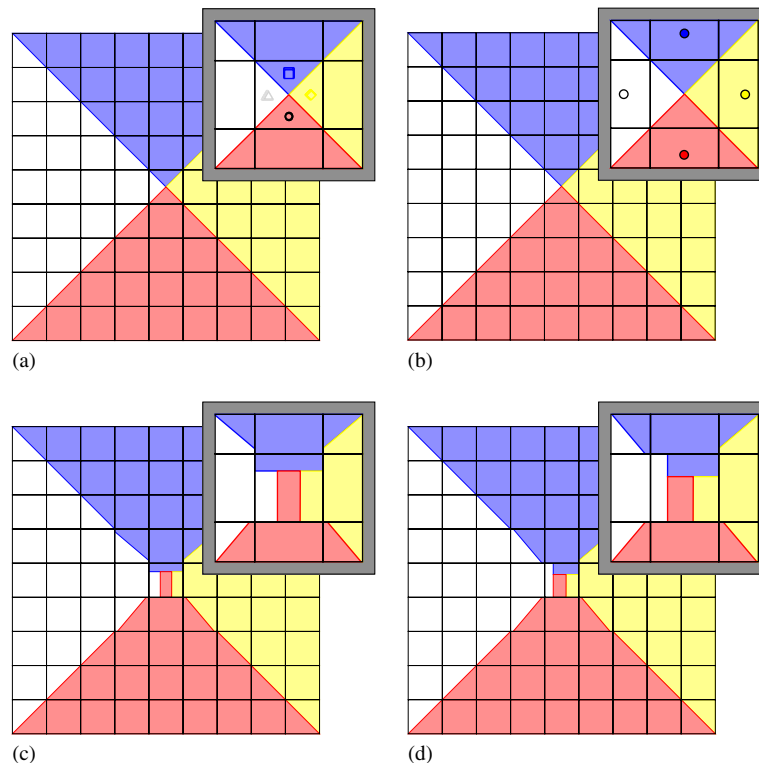


Figure 8. Four-material interface reconstruction using: (a) particles and power diagrams; (b) approximate centroids and power diagrams; and (c), (d) Youngs' method with two different material orderings. The insets show the four-material cell at the centre of the mesh. The converged particle locations for the centre cell are also shown in (a). The converged approximate centres of mass for the centre cell are shown in (b).

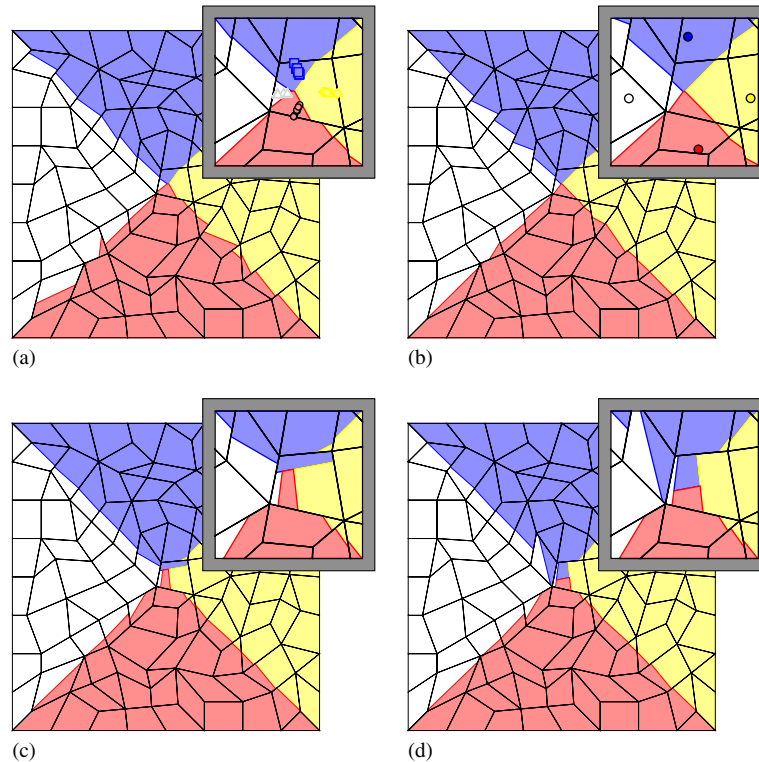


Figure 9. Four-material interface reconstruction on an unstructured grid using: (a) particles and power diagrams; (b) approximate centroids and power diagrams; and (c), (d) Youngs' method with two different material orderings. The insets show the four material cell at the centre of the mesh with the particles for the centre cell in (a) and the approximate centres of mass for the centre cell in (b).

the multi-material centre cell, the power diagram-based reconstructions in Figures 9(a) and (b) are more representative of the structure of the true configuration. In Figures 9(c) and (d), the material order dependency provides significantly different reconstructions for the four-material centre cell. This and the structured example in Figure 8 demonstrate the strength of our method in reconstructing multi-material cells. In two-material cells, Youngs' and related methods do quite well, as is shown in Figures 9(c) and (d) in the regions away from the centre. An effective strategy is to *use existing methods for two-material cells, and a power diagram-based method for cells containing three or more materials*. In Figures 8 and 9, the power diagram reconstruction was used on the entire mesh, which would not typically be done in practice.

Figure 10 shows the reconstruction of a filament-type structure that is not aligned with the grid. The filament is preserved with all the three-material cells showing the proper material positions. The power diagram reconstructions shown in Figures 10(c) and (d) do not reverse the location of the materials relative to the filament as does Youngs' reconstruction with the wrong material ordering as shown in Figure 10(b). The power diagram-based method fails to exactly reproduce a straight line in both Figures 10(c) and (d). The reconstruction in the cells at the top and bottom of the grid in Figures 10(c) and (d) could be improved with better boundary treatment. No

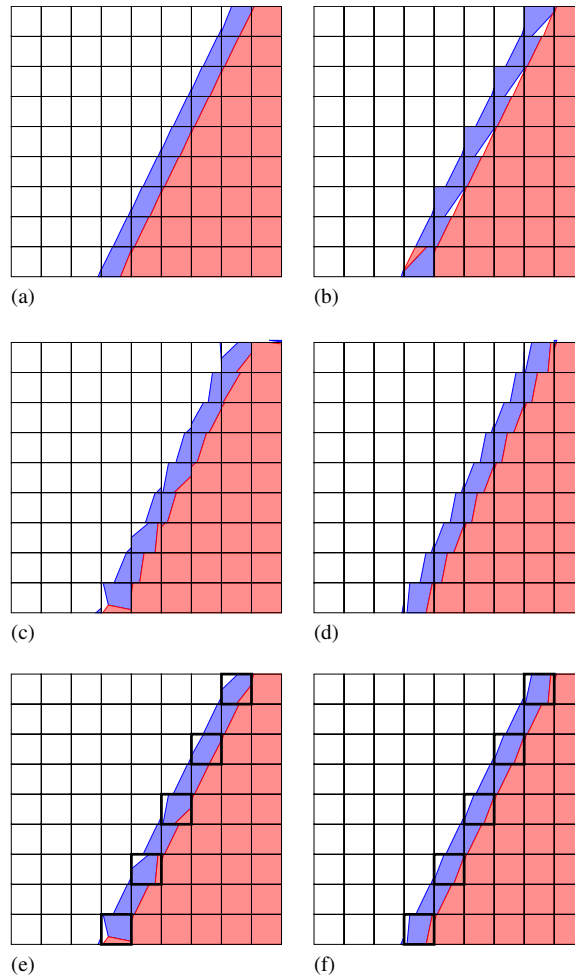


Figure 10. Multi-material interface reconstruction for a filament: Youngs' reconstruction with (a) correct material ordering and (b) incorrect material ordering; (c) particles and power diagram reconstruction in all cells; (d) approximate centroids and power diagram-based reconstruction in all cells; (e) mixed approach using Youngs' method for two-material cells and particles and power diagram reconstruction in three or more material cells; and (f) mixed approach using Youngs' method and approximate centroids and power diagram reconstruction. The three-material cells using the power diagram-based reconstructions are in bold in (e) and (f). Notice the presence of the white material on the right-hand side of the filament in Youngs' reconstruction in (b). This can lead to 'flotsam and jetsam'.

special treatment was used at the boundary for either the particle method or approximate centroid method. With the correct material ordering, Youngs' method produces a very good reconstruction as shown in Figure 10(a). This motivates the use of a mixed method using Youngs' method for two-material cells and a power diagram-based reconstruction for three or more material cells. These reconstructions are shown in Figures 10(e) and (f). The multi-material cells, shown with bold outlines, used a power diagram-based method while the rest of the interface was reconstructed

using Youngs' method. These reconstructions do not have the material ordering problems as would Youngs' method used on all mesh cells.

6. CONCLUSIONS

We have developed a first-order method for material order-independent reconstruction of multi-material cells. It is based on the determination of the relative locations of each material in a cell and then reconstructing the interface using a power diagram, that once clipped to the mesh cell, matches the desired volume fractions. The method does not assume a topology for the material regions, i.e. a layer structure or triple point configuration. Furthermore, all of the material regions created will be convex.

The performance of the particle model is good. The average particle position converges to within a few per cent of its converged position in under 20 time steps of the attraction–repulsion model. Furthermore, the particles do not show much sensitivity to the random initial conditions used.

The approximate centre of mass approach typically gives better results, and relies on no externally supplied parameters. In addition, while being highly problem dependent, it can be around 50 times faster than the particle method. Neither the particle model nor approximate centre of mass method when combined with a power diagram-based reconstruction exactly reproduces a straight line, indicating both methods are only first-order accurate. While the potential for the particles to detect the presence of a thin filament within the mesh through the determination of multiple particle clusters is intriguing, we recommend the approximate centre of mass approach, particularly when subcell details such as thin filaments do not have to be resolved.

The power diagram-based interface reconstruction maintains the relative location of the materials within the cell and may be useful in other methods utilizing approximate material location information.

As two-material reconstruction algorithms are reliable and well established, this method would be most applicable for reconstruction of only cells containing more than two materials. This reconstruction method is currently being implemented in multi-material flow codes to further understand the combined errors of advection and interface reconstruction.

ACKNOWLEDGEMENTS

The authors would like to thank B. K. Swartz and M. J. Shashkov for their helpful discussions. This work was performed under the auspices of the National Nuclear Security Administration of the US Department of Energy at Los Alamos National Laboratory under Contract No. DE-AC52-06NA25396 and supported by the DOE Advanced Simulation and Computing (ASC) program.

REFERENCES

1. Benson DJ. Computational methods in Lagrangian and Eulerian hydrocodes. *Computer Methods in Applied Mechanics and Engineering* 1992; **99**:235–394.
2. Benson DJ. Volume of fluid interface reconstruction methods for multi-material problems. *Applied Mechanics Review* 2002; **55**(2):151–165.
3. Rider WJ, Kothe DB. Reconstructing volume tracking. *Journal of Computational Physics* 1998; **141**:112–152.
4. Scardovelli R, Zaleski S. Direct numerical simulation of free-surface and interfacial flow. *Annual Review of Fluid Mechanics* 1999; **31**:567–603.

5. Hirt CW, Nichols BD. Volume of fluid (VOF) method for the dynamics of free boundaries. *Journal of Computational Physics* 1981; **39**:201–225.
6. Noh WF, Woodward P. SLIC (Simple line interface calculation). In *Lecture Notes in Physics*, vol. 59, van de Vooren AI, Zandbergen PJ (eds). Springer: Berlin, 1976; 330–340.
7. Youngs DL. Time dependent multi-material flow with large fluid distortion. In *Numerical Methods for Fluid Dynamics*, Morton KW, Baines MJ (eds). Academic Press: New York, 1982; 273–285.
8. Youngs DL. An interface tracking method for a 3d Eulerian hydrodynamics code. *Technical Report AWE/44/92/35*, AWRE Design and Math Division, 1984.
9. Barth TJ. Aspects of unstructured grids and finite-volume solvers for Euler and Navier–Stokes equations. *VKI/NASA/AGARD Special Course on Unstructured Grid Methods for Advection Dominated Flows AGARD Publication R-787*, 1995.
10. Pilliod Jr JE, Puckett EG. Second-order accurate volume-of-fluid algorithms for tracking material interfaces. *Journal of Computational Physics* 2004; **199**:465–502.
11. Garimella RV, Dyadechko V, Swartz BK, Shaskov MJ. Interface reconstruction in multi-fluid, multi-phase flow simulations. *Proceedings of the 14th International Meshing Roundtable*, San Diego, CA. Springer: Berlin, 2005; 19–32.
12. Swartz BK. The second-order sharpening of blurred smooth borders. *Mathematics of Computation* 1989; **52**(186):675–714.
13. Choi BY, Bussmann M. A piecewise linear approach to volume tracking a triple point. *International Journal for Numerical Methods in Fluids* 2007; **53**:1005–1018.
14. Dyadechko V, Shashkov MJ. Multi-material interface reconstruction from the moment data. *Technical Report LA-UR-06-5846*, Los Alamos National Laboratory, Los Alamos, NM, 2006.
15. Mosso S, Clancy S. A geometrically derived priority system for Young’s interface reconstruction. *Technical Report LA-CP-95-0081*, Los Alamos National Laboratory, Los Alamos, NM, 1995.
16. Benson DJ. Eulerian finite element methods for the micromechanics of heterogeneous materials: dynamic prioritization of material interfaces. *Computer Methods in Applied Mechanics and Engineering* 1998; **151**: 343–360.
17. Allen MP, Tildesley DJ. *Computer Simulation of Liquids*. Oxford University Press: New York, 1987.
18. Haile JM. *Molecular Dynamics Simulation: Elementary Methods*. Wiley: New York, 1992.
19. Benz W. Smooth particle hydrodynamics: a review. *The Numerical Modelling of Nonlinear Stellar Pulsations: Problems and Prospect*. NATO ASI Series. Series C, Mathematical and Physical Sciences, vol. 302. Kluwer Academic Publishers: Dordrecht, 1989; 269–287.
20. Gingold RA, Monaghan JJ. Smoothed particle hydrodynamics—theory and application to non-spherical stars. *Monthly Notices of the Royal Astronomical Society* 1977; **181**:375–389.
21. Schofield SP, Garimella RV, Francois MM, Loubere R. Multi-material interface reconstruction using particles and power diagrams. *Technical Report LA-UR-06-8740*, Los Alamos National Laboratory, Los Alamos, NM, 2006.
22. Jain AK, Murty MN, Flynn PJ. Data clustering: a review. *ACM Computing Surveys* 1999; **31**(3):264–323.
23. Puckett EG, Saltzman JS. A 3d adaptive mesh refinement algorithm for multimaterial gas dynamics. *Physica D* 1992; **60**:84–93.
24. Aurenhammer F. Power diagrams: properties, algorithms and applications. *SIAM Journal on Computing* 1987; **16**(1):78–96.
25. Imai H, Iri M, Murota K. Voronoi diagram in the Laguerre geometry and its applications. *SIAM Journal on Computing* 1985; **14**(1):93–105.
26. Edelsbrunner H, Shah NR. Incremental topological flipping works for regular triangulations. *Algorithmica* 1996; **15**:223–241.
27. Anton F, Mioc D. On the conversion of ordinary Voronoi diagrams into Laguerre diagrams. *The Eleventh Canadian Conference on Computational Geometry*, Vancouver, Canada, August 1999; 150–153.