

A Comparison of Performance between
Reconstruction and Advection Algorithms
for Volume-of-Fluid Methods

by

Dominic Kedelty

A Thesis Presented in Partial Fulfillment
of the Requirements for the Degree
Master of Science

Approved April 2015 by the
Graduate Supervisory Committee:

Marcus Herrmann, Chair
Huei-Ping Huang
Kangping Chen

ARIZONA STATE UNIVERSITY

May 2015

ABSTRACT

The Volume-of-Fluid method is a popular method for interface tracking in Multiphase applications within Computational Fluid Dynamics. To date there exists several algorithms for reconstruction of a geometric interface surface. Of these are the Finite Difference algorithm, Least Squares Volume-of-Fluid Interface Reconstruction Algorithm, LVIRA, and the Efficient Least Squares Volume-of-Fluid Interface Reconstruction Algorithm, ELVIRA. Along with these geometric interface reconstruction algorithms, there exist several volume-of-fluid transportation algorithms. This paper will discuss two operator-splitting advection algorithms and an unsplit advection algorithm. Using these three interface reconstruction algorithms, and three advection algorithms, a comparison will be drawn to see how different combinations of these algorithms perform with respect to accuracy as well as computational expense.

TABLE OF CONTENTS

	Page
LIST OF TABLES	iv
LIST OF FIGURES	v
CHAPTER	
1 INTRODUCTION	1
2 BACKGROUND	3
Definition of the Volume Fraction	3
Piecewise Linear Interface Calculation	4
Transportation of the Volume Fraction.....	5
3 INTERFACE RECONSTRUCTION	7
Finite Difference Formulation	8
Least-Squares Volume-of-Fluid Interface Reconstruction Algorithm	9
Efficient Least-Squares VoF Interface Reconstruction Algorithm	11
4 FLUID ADVECTION	12
Operator-Split Non-Rotational Advection	13
Operator-Split Rotational Advection.....	14
Conservative Unsplit Geometric Advection	16
5 RESULTS: TEST CASES	18
Reconstruction of a Randomly Placed Circle	18
Zalesak's Disk Test Case Definition	20
Zalesak's Disk Results: Operator-Split Non-Rotational Advection.....	23

CHAPTER	Page
Zalesak’s Disk Results: Unsplit Advection.....	26
Zalesak’s Disk Results: Advection Timing Comparison.....	29
Circle in a Deformation Field Definition	31
Circle in a Deformation Field Results: Non-Rotational Advection	32
Circle in a Deformation Field Results: Rotational Advection.....	39
Circle in a Deformation Field Results: Unsplit Advection.....	46
Circle in a Deformation Field Results: Advection Timing Comparison	53
Sphere in a Deformation Field Definition	54
Sphere in a Deformation Field Results: Non-Rotational Advection	55
Sphere in a Deformation Field Results: Rotational Advection	61
Sphere in a Deformation Field Results: Unsplit Advection	67
Sphere in a Deformation Field Results: Advection Timing Comparison	74
6 CONCLUSION	75
REFERENCES.....	80

LIST OF TABLES

Table		Page
1.	Report of L_1 Error Norms in Randomly Placed Circle Test Case	19
2.	Zalesak's Disk Operator-Split Non-Rotational Shape Errors	24
3.	Time per Iteration Comparison of Three Normal Schemes	25
4.	Zalesak's Disk Unsplit Shape Errors	26
5.	Time per Iteration Comparison of Three Normal Schemes	28
6.	Time per Iteration Comparison of Three Advection Schemes	30
7.	DC: Operator-Split Non-Rotational Shape Errors.....	37
8.	Time per Iteration Comparison of Three Normal Schemes	38
9.	DC: Operator-Split Rotational Shape Errors	44
10.	Time per Iteration Comparison of Three Normal Schemes	45
11.	DC: Unsplit Shape Errors.....	51
12.	Time per Iteration Comparison of Three Normal Schemes	52
13.	Time per Iteration Comparison of Three Advection Schemes	53
14.	DS: Operator-Split Non-Rotational Shape Errors	59
15.	Time per Iteration Comparison of Three Normal Schemes	60
16.	DS: Operator-Split Rotational Shape Errors	65
17.	Time per Iteration Comparison of Three Normal Schemes	66
18.	DS: Unsplit Shape Errors	72
19.	Time per Iteration Comparison of Three Normal Schemes	73
20.	Time per Iteration Comparison of Three Advection Schemes	74

LIST OF FIGURES

Figure		Page
1.	Interface Between Two Fluids	3
2.	Piecewise Linear Interface Calculation	5
3.	Exact Interface to Volume Fraction	7
4.	Finite Difference Normal Depiction.....	8
5.	LVIRA and ELVIRA Normal Depiction	9
6.	Example of Operator-Split Non-Rotational Advection Fluxing	14
7.	Example of Operator-Split Rotational Advection Fluxing	15
8.	Example of Geometric Unsplit Advection Fluxing	16
9.	Convergence plot of L_1 Errors in Approximating Normal Directions	19
10.	Zalesak's Disk Geometry Definition	20
11.	Results of the Zalesak's Disk Test Case using Non-Rotational Advection	22
12.	Close-up view of the top right corner of the notch in Zalesak's disk	23
13.	Shape Error Plot of Zalesak's Disk using Non-Rotational Advection.....	24
14.	Bar Graph Timing Comparison of the Three Normal Approximations	25
15.	Shape Error Plot of Zalesak's Disk using Unsplit Advection	26
16.	Bar Graph Timing Comparison of the Three Normal Approximations.....	27
17.	Results of the Zalesak's Disk Test Case using Unsplit Advection	28
18.	Bar Graph Timing Comparison of the Three Advection Schemes	30
19.	Visual of the Circle in a Deformation Field Test Case	31
20.	Results of the DC Test Case using Non-Rotational Advection – T=4	33

Figure	Page
21. Shape Error vs. Simulation Time – Non-Rotational Advection	34
22. Results of the DC Test Case using Non-Rotational Advection – T=8.....	36
23. Shape Error Plot of DC Test Case using Non-Rotational Advection	37
24. Bar Graph Timing Comparison of the Three Normal Approximations	38
25. Results of the DC Test Case using Rotational Advection – T=4	40
26. Shape Error vs. Simulation Time – Rotational Advection	41
27. Results of the DC Test Case using Rotational Advection – T=8	43
28. Shape Error Plot of DC Test Case using Rotational Advection.....	44
29. Bar Graph Timing Comparison of the Three Normal Approximations	45
30. Results of the DC Test Case using Unsplit Advection – T=4	47
31. Shape Error vs. Simulation Time – Rotational Advection	48
32. Results of the DC Test Case using Unsplit Advection – T=8	49
33. Shape Error Plot of DC Test Case using Unsplit Advection.....	51
34. Bar Graph Timing Comparison of the Three Normal Approximations	52
35. Bar Graph Timing Comparison of the Three Advection Schemes	53
36. Visual of the Sphere in a Deformation Field Test Case	54
37. Results of the DS Test Case using Non-Rotational Advection – T=1.5 ...	56
38. Shape Error vs. Simulation Time – Non-Rotational Advection.....	57
39. Results of the DS Test Case using Non-Rotational Advection – T=3	58
40. Shape Error Plot of DS Test Case using Non-Rotational Advection.....	59

Figure	Page
41. Bar Graph Timing Comparison of the Three Normal Approximations	60
42. Results of the DS Test Case using Rotational Advection – T=1.5	62
43. Shape Error vs. Simulation Time – Rotational Advection	63
44. Results of the DS Test Case using Rotational Advection – T=3.....	64
45. Shape Error Plot of DS Test Case using Rotational Advection	65
46. Bar Graph Timing Comparison of the Three Normal Approximations.....	66
47. Results of the DS Test Case using Unsplit Advection – T=1.5.....	68
48. Shape Error vs. Simulation Time – Unsplit Advection	69
49. Results of the DS Test Case using Unsplit Advection – T=3.....	70
50. Shape Error Plot of DS Test Case using Unsplit Advection	72
51. Bar Graph Timing Comparison of the Three Normal Approximations	73
52. Bar Graph Timing Comparison of the Three Advection Schemes	74
53. Shape Error vs Average Time per Timestep for DC Test Case	77

CHAPTER 1

INTRODUCTION

The Volume-of-Fluid algorithm is a popular interface-capturing method for the tracking of a free surface in numerical simulation of a fluid dynamics problem. Along the Volume-of-Fluid algorithm calculates the transportation of a given interface geometry in a known velocity field. However, when coupled with a numerical flow solver like the discrete Navier-Stokes equations, the motion of two-phase flows can be predicted. For accurate prediction and understanding of interfacial flows, an accurate Volume-of-Fluid interface-capturing scheme is needed. Many different re-formulations of the Volume-of-Fluid method exist.

In this paper, a comparison will be made of different schemes for segments of the Volume-of-Fluid algorithm. Firstly, the scheme for determining the interface normal within a given cell will be examined. Several methods exist; examples are Parker & Young's method (1992), the finite difference (Li, 1995), the Least-Squares Volume-of-Fluid Interface Reconstruction Algorithm, LVIRA (Puckett, 1991), and Efficient Least-Squares Volume of Fluid Method, ELVIRA (Pilliod & Puckett, 2004). These interface normal calculations have been devised as a way to best calculate the straight-line interface normal using only information of the volume of a geometry. Since the true interface in a given cell is not always straight, it is difficult to define a correct normal. This paper focuses on the LVIRA, and ELVIRA methods for determining the normal interface direction in the reconstruction step of the Volume-of-Fluid. The finite difference normal is also compared as a baseline as it is one of the most widely used normal calculations for the Volume-of-Fluid.

Second, the advection scheme for transportation of the interface is compared. Many deviations exist for the purpose of the fastest, most accurate, and robust transportation of the interface. A comparison will be drawn using the following transportation algorithms; advection of the reconstructed interface taking the flux face velocity as constant over an entire cell (Scardovelli & Zaleski, 1999), advection of the reconstructed interface interpolating the velocity across a cell volume that allows stretching and compression of the interface (Gueyffier *et al.* 1999), and finally a second order conservative unsplit algorithm (Owkes & Desjardins, 2014).

In order to compare the different types of interface reconstruction and interface transportation schemes, widely used test cases will be used to determine the error accuracy and relative speed between the different volume-of-fluid advection and reconstruction schemes. The test cases used in this work are the Zalesak's disk, the circular column in a deformation field test case, and a three-dimensional sphere in a deformation field test case. These test cases were chosen because of they are widely used in literature and have documented information found in Owkes & Desjardins (2014) as well as Pilliod & Puckett (2004).

This motivation behind this work is to evaluate the impact of the different interface normal calculations and transportation algorithms on the Volume-of-Fluid method. The paper is organized as follows. Chapter 2 gives a detailed description of the processes in the Volume-of-Fluid method. Chapter 3 explains the different types of normal calculations that are used in this paper. Chapter 4 gives an explanation of the different transportation methods. Chapter 5 contains the results for the following test cases evaluated within this work.

CHAPTER 2

BACKGROUND

Definition of the Volume Fraction

Beginning with a problem of two immiscible fluids separated by a free surface inside a spatial domain. One can differentiate between the two fluids by defining one as a dark fluid, and the other as a light fluid. The spatial domain is decomposed into finite sized control volumes. Instead of explicitly tracking the free surface, the separation of fluids can be captured by tracking the ratio of the dark fluid volume to total volume within each cell. The ratio of dark fluid to cell volume is called the liquid volume fraction, C . In some literature, such as (Gueyffier *et al.* 1999), the ratio is called the color function. By storing only the liquid volume fraction, information about the exact interface between the dark and light fluid is lost.




Figure 1 – An interface defined within a spatial domain. The grid shows cells that have been discretized as finite sized control volumes. Cells are colored according to their volume fraction. A dark grey cell denotes a volume fraction of one, while a white cell denotes a volume fraction of zero.

Since information of the exact interface is lost, the initial portion of the Volume-of-Fluid algorithm is dedicated to using the volume fraction to approximately reconstruct an interface. This reconstruction is accomplished geometrically using simple elements that geometrically represent the volume within each individual control volume. A number of interface reconstruction schemes are available such as the Simple Line Interface Calculation, SLIC (Noh & Woodward, 1976), and the Piece-wise Line Interface Reconstruction, PLIC (Gueyffier *et al.* 1999). An attempt has even been made for a Parabolic Interface Reconstruction scheme, PROST (Renardy & Renardy, 2002). However due to issues with robustness, the parabolic scheme is not widely used.

Piecewise Linear Interface Calculation

In this work, the PLIC reconstruction scheme is used for the reconstruction of the interface between immiscible fluids. The PLIC reconstruction scheme is the following. Linear elements are used to describe the volume fraction within each control volume. The interface between fluids is then defined using 1st order linear segments in each cell.

$$x_1 m_1 + x_2 m_2 + x_3 m_3 = \alpha \quad (1)$$

Each cell contains its own linear segment, and therefore the linear segments across multiple cells are piece-wise discontinuous. In reconstruction, the linear segment defined by equation 1 must be identified. This requires an approximation for the interface normal, \mathbf{m} , as well as the perpendicular distance from the cell origin, α . Calculations for determining the interface normal will be discussed in Chapter 3.

After the appropriate normal is calculated, the perpendicular distance from the cell origin must be determined. This can be accomplished by constraining equation 1

such that the volume, under the plane is equal the volume fraction of the cell, using equation 2.

$$C = \frac{1}{6m_1m_2m_3} \left[\alpha^3 - \sum_{i=1}^3 H(\alpha - m_i)(\alpha - m_i)^3 - \sum_{i=1}^3 H(\alpha - \alpha_{max} + m_i) \right] \quad (2)$$

$$\alpha_{max} = m_1 + m_2 + m_3$$

The heaviside function is made use of and is denoted by H. It should be noted that equation 2 assumes an equidistant cell size of 1. This is because it is the volume fraction that is of interest, not the dark liquid volume. The interface segment that is reconstructed within a cell volume is depicted in figure 2.




Figure 2 – Piecewise linear interface reconstruction of one cell within the computational domain.

Transportation of the Volume Fraction

After reconstruction of the interface, the transportation of the dark liquid is required. The PLIC reconstruction of the interface makes it difficult to perform advection using an advection equation on the reconstructed interface since it is piecewise discontinuous and defined locally for each cell. Therefore transportation is accomplished

through the geometric fluxing of the volume fraction through cell faces. This requires that the velocities within the domain be defined normal to cell faces.

The fluxes are determined at each face of the control volumes within the spatial domain. When the fluxes are determined, the volume fractions are updated. This is accomplished by subtracting the fluxed volume fraction, and placing it into the adjacent cell volume. Adjacent meaning, the cell that shares the fluxing face. By using this direct fluxing operation, the Volume-of-Fluid method is discretely conservative. Three different methods for the transportation of the volume fraction are examined in this work. Detailed instruction of the fluxing operation is given in Chapter 4.

CHAPTER 3

INTERFACE NORMAL CALCULATION

We start with the interface that passes through the center cell of a 3-by-3-cell stencil. A 3-by-3-by-3-cell stencil is used in three-dimensions, but 2D formulations will be described here since they are easily extendible to 3D. Figure 3, shows an interface example that passes through the center cell of the stencil.




Figure 3 – Left: Exact interface within a 3-by-3-cell stencil. Right: Volume fraction decomposition of the exact interface.

The volume fraction, is the only known information. Figure 3b, shows the volume fractions, which are approximate for this example. The interface is then calculated in the next sections by the following; Finite Difference, LVIRA, and ELVIRA.

Finite Difference Formulation

The finite difference formulation is due to Li (1995). In this formulation, the normal of the interface is approximated by a finite difference approximation of the non-unit normal following,

$$\mathbf{m}^h = -\nabla^h C \quad (3)$$

In order to define a stable normal for the interface normal reconstruction, gradients are calculated at cell corners and averaged back to the cell centers such that.

$$\mathbf{m}_{i,j} = \frac{1}{4} \left(\mathbf{m}_{i+\frac{1}{2},j-\frac{1}{2}} + \mathbf{m}_{i-\frac{1}{2},j-\frac{1}{2}} + \mathbf{m}_{i+\frac{1}{2},j+\frac{1}{2}} + \mathbf{m}_{i-\frac{1}{2},j+\frac{1}{2}} \right) \quad (4)$$

In three-dimensions the interface normal is the average of the eight corners of a cubic cell.

$$\begin{aligned} \mathbf{m}_{i,j,k} = \frac{1}{8} \left(\mathbf{m}_{i+\frac{1}{2},j-\frac{1}{2},k+\frac{1}{2}} + \mathbf{m}_{i-\frac{1}{2},j-\frac{1}{2},k+\frac{1}{2}} + \mathbf{m}_{i+\frac{1}{2},j+\frac{1}{2},k+\frac{1}{2}} + \mathbf{m}_{i-\frac{1}{2},j+\frac{1}{2},k+\frac{1}{2}} + \mathbf{m}_{i+\frac{1}{2},j-\frac{1}{2},k-\frac{1}{2}} \right. \\ \left. + \mathbf{m}_{i-\frac{1}{2},j-\frac{1}{2},k-\frac{1}{2}} + \mathbf{m}_{i+\frac{1}{2},j+\frac{1}{2},k-\frac{1}{2}} + \mathbf{m}_{i-\frac{1}{2},j+\frac{1}{2},k-\frac{1}{2}} \right) \quad (5) \end{aligned}$$

Figure 4, shows visually the calculation of the finite difference normal approximation.

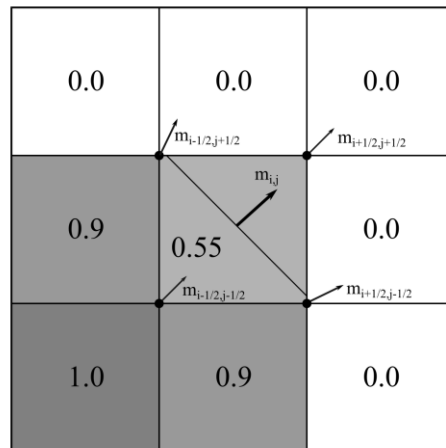


Figure 4 – Visual depiction of how the finite difference normal is obtained. The four vectors at cell vertices are averaged to the cell center to determine the interface of the line segment.

Least Squares Volume-of-Fluid Interface Reconstruction Algorithm

The Least-Squares Volume-of-Fluid Interface Reconstruction Algorithm is due to Puckett (1991). In this approach a minimum error metric is used to find the best normal approximation. Puckett approximates the best normal by drawing an interface as a straight line that passes through the entire 3-by-3-cell stencil, figure 5.




Figure 5 – LVIRA and ELVIRA normal approximation depictions. Left: true volume fraction known from the exact interface. Right: approximated volume fractions of every cell in the 3-by-3-cell stencil when the line segment is drawn through the entire 3-by-3-cell stencil.

The interface in figure 5, is constructed such that the volume fraction given under the interface reconstruction in the center cell is the true value of C . Under this constraint, an error metric can be defined as the straight-line interface that minimizes the difference between reconstructed volume fractions and true volume fractions in the surrounding cells.

$$||\mathbf{E}(\mathbf{m})|| = \left(\sum_{k,l=-1}^1 (\tilde{C}_{i+k,j+l}(\mathbf{m}) - C_{i+k,j+l})^2 \right)^{1/2} \quad (6)$$

This then becomes a non-linear least squares optimization problem. To solve the non-linear least squares problem imposed by LVIRA, the Levenberg-Marquardt gradient decent algorithm was chosen to quickly and accurately find the minimum value in the error function,

$$[\mathbf{J}^T] + \lambda \mathbf{diag}(\mathbf{J}^T \mathbf{J}) \boldsymbol{\delta} = \mathbf{J}^T \mathbf{E} \quad (7)$$

Where the following are defined as,

$$\mathbf{J} = \frac{\partial \mathbf{E}}{\partial \mathbf{m}} = \frac{\mathbf{E}(\theta) - \mathbf{E}(\theta + \Delta \theta)}{\Delta \theta}$$

$\lambda = \text{user defined damping factor}$

$$\boldsymbol{\delta} = [\beta_1, \beta_2, \dots, \beta_n]^T$$

On the right hand side of equation 7 is the vector, \mathbf{E} , containing the error between approximated volume fraction and true volume fraction in each cell. In two-dimensions, equation 6 is a function of two parameters, m_1 and m_2 . To minimize the number of parameters affecting the error function, the non-unit normal vector is converted to an angle, θ . This reduces the number of parameters from m_1 and m_2 to a single parameter, θ . In three-dimensions, the number of parameters is reduced from m_1, m_2, m_3 to θ, ϕ .

The Levenberg-Marquardt method is a gradient decent algorithm for finding the local minima within the vicinity of an initial guess. For the comparisons of the different normal calculation schemes, the initial guess that was chosen was the ELVIRA normal. The reason ELVIRA is chosen as the initial guess is that the lowest error in equation 6 is of interest for this work. An initial guess using the finite difference formulation has shown higher errors than when using ELVIRA. This is the result of multiple local minima limiting the Levenberg-Marquardt from finding the global minimum.

Efficient Least Squares Volume-of-Fluid Interface Reconstruction Algorithm

The Efficient Least-Squares Volume-of-Fluid Interface Reconstruction Algorithm is due to Pilliod & Puckett (2004). In this method the same error metric as used in the LVIRA method is used. To reduce the computational workload, instead of iterating to find a minimum solution, the ELVIRA method chooses between several candidates of \mathbf{m} . The candidates that are selected are the forward, central, and backward differences of the column sums in each of the principle axis directions. In two dimensions, the total number of candidates is 6, and in three dimensions the number of candidates is 9. The candidates in two-dimensions are the following,

$$m_b^x = \sum_{l=-1}^1 f_{i,j+l,k} - f_{i-1,j+l,k} \quad (8)$$

$$m_c^x = \frac{1}{2} \sum_{l=-1}^1 f_{i+1,j+l,k} - f_{i-1,j+l,k} \quad (9)$$

$$m_f^x = \sum_{l=-1}^1 f_{i+1,j+l,k} - f_{i,j+l,k} \quad (10)$$

This is similar for m_b^y , m_c^y , and m_f^y . From the criterion, the appropriate normal is chosen such that it has the smallest error in the error criterion shown in equation 6.

$$\tilde{\mathbf{n}} = \min\{E(m_b^x, m_c^x, m_f^x, m_b^y, m_c^y, m_f^y)\} \quad (11)$$

Pilliod and Puckett (2004), claim that for straight line interfaces both the ELVIRA, and LVIRA method produce the exact normal regardless of the orientation. However for curved line interfaces the LVIRA should produce better results at a higher computational cost.

CHAPTER 4
FLUID ADVECTION

Common in practice, the fluxing operation of the advection algorithm is broken into separate operator-splitting steps. This is because it significantly reduces the complexity of the advection algorithm. Operator-split transportation methods reduce the complexity by reducing the amount of fluxing that is possible during a time-step. Given the advection equation.

$$\frac{\partial C}{\partial t} + \mathbf{u} \cdot \nabla C = 0 \quad (12)$$

The operator-split transportation scheme becomes

$$\frac{\partial C}{\partial t} + u_x \frac{\partial C}{\partial x} = 0 \quad (13)$$

$$\frac{\partial C}{\partial t} + u_y \frac{\partial C}{\partial y} = 0 \quad (14)$$

$$\frac{\partial C}{\partial t} + u_z \frac{\partial C}{\partial z} = 0 \quad (15)$$

In each operator-split step, the interface is transported along each of the principle directions independently. By doing this the fluxes through two faces at most must be calculated for each operator-split step. This however has the drawback of transporting the volume fraction during non-divergence free fluxing steps. Transportation errors due to these independent fluxing steps are known as splitting errors. These splitting errors lead to unwanted elements in the domain known as float-sum in which the volume fraction contained within cells can fall below 0 and in excess of 1.

Operator-Split Non-Rotational Advection

The non-rotational operator-split method is due to Scardovelli & Zaleski (1999). In non-rotational operator-split transportation, only the translation of the interface is considered. The translation of the interface can be described as follows. Starting with equation 1 for a planar interface that is defined within a cell volume, with n denoting the n^{th} time step,

$$x_1^n m_1^n + x_2^n m_2^n + x_3^n m_3^n = \alpha^n \quad (16)$$

Denoting the fluxing face velocity, U_h , the position of each point on the interface can be calculated from the displacement during the time-step.

$$x_1^{n+1} = x_1^n + U_h \tau \quad (17)$$

Substituting (17) into (16), the equation for the interface after transportation from the fluid flow becomes equation 18.

$$x_1^{n+1} m_1^n + x_2^n m_2^n + x_3^n m_3^n = \alpha^n + (U_h \tau) m_1^n \quad (18)$$

For operator-splitting, the translation in other principal directions for this step are ignored and the equation for planar interface after transportation along the x_1 principal direction becomes

$$x_1^* m_1^* + x_2^* m_2^* + x_3^* m_3^* = \alpha^* \quad (19)$$

$$\alpha^* = \alpha^n + U_h \tau m_1^n$$

The star values denote the equation between operator-splitting steps, since this fluxing step must be done for each of the principal directions. Observation of equation 19 shows that the transportation of the interface results in a repositioning of the origin of the coordinate system. Fluxes are then calculated as the volume under the line that has passed through the flux face, figure 6.




Figure 6 – Operator-split non-rotational advection. Only the translation of the interface is taken into account. All points on the line are translated using the flux face velocity over a time-step, Δt .

Operator-Split Rotational Advection

Rotational operator-split transportation is due to Gueyffier *et al.* (1999). In this version of the operator split fluxing, a more complete method of interface transportation is considered. Instead of calculating fluxes through a face by taking the face velocity as constant over the entire cell, this scheme calculates the Lagrangian motion of the interface by interpolation using opposite face velocities, U_0 and U_h . The interface equation after advection can be calculated similar to the non-rotational transportation scheme. By assuming a unit sized computational cell, the x_1 component of the velocity within the cell can be interpolated across the cell.

$$u_1(x_1) = U_0 \left(1 - \frac{x_1}{h}\right) + U_h \frac{x_1}{h} \quad (20)$$

Each point on the interface at the beginning of the time-step is transported using the above equation assuming that the face velocities remain constant during the entire advection step. The x_1 coordinate of each point on the initial interface is repositioned.

$$x_1^* = x_1^n + u_1(x_1^n)\tau = \left[1 + \left(\frac{U_h - U_0}{h}\right)\tau\right] x_1^n + U_0\tau \quad (21)$$

Substituting equation 21 into equation 22.

$$x_1^n = \frac{x_1^* - U_0 \tau}{1 + \left(\frac{U_h - U_0}{h}\right) \tau} \quad (22)$$

The equation for the planar surface after the operator split step is calculated by substituting equation 22 into equation 16.

$$m_1^n \left[\frac{x_1^* - U_0 \tau}{1 + \left(\frac{U_h - U_0}{h}\right) \tau} \right] + m_2^n x_2^n + m_3^n x_3^n = \alpha^n \quad (23)$$

The remaining x_2 and x_3 coordinates remain unchanged during this advection step along the x_1 principle direction step, and thus for this step in the operator splitting advection, the new equation for the interface becomes

$$m_1^* x_1^* + m_2^* x_2^* + m_3^* x_3^* = \alpha^* \quad (24)$$

$$m_1^* = \frac{m_1^n}{1 + \left(\frac{U_h - U_0}{h}\right) \tau}$$

$$\alpha^* = \alpha^n + \frac{m_1^n U_0 \tau}{1 + \left(\frac{U_h - U_0}{h}\right) \tau}$$

All other star values remain equal to their n^{th} values during advection along the x_1 direction. Figure 7, shows the flux volume after the interface has been transported along the x_1 direction. This is repeated for the remaining principle directions.




Figure 7 – Operator-split rotational advection. The interface segment is propagated using opposite face velocities interpolated across the control volume. This allows stretching and compressing.

Unsplit Advection

The unsplit advection scheme tested in this work is the conservative three-dimensional unsplit geometric transport by Owkes and Desjardins (2014). This method is by far the most complicated method. For an in-depth understanding of the method, refer to Owkes and Desjardins (2014).




Figure 8 – Unsplit geometric fluxing used in the 2nd order unsplit transportation method. The fluxed volume fraction is the volume under the PLIC surface that is contained within the flux volume calculated by back-tracing cell vertices in time.

In this method the flux volume across a cell face is approximated using 3D fluxing polyhedra calculated by back-tracing cell vertices in time. A 3D version of this method is shown in figure 8. Using cell vertices to construct fluxing volumes is done to ensure conservation of mass during fluxing. This way flux vertices never overlap causing double fluxing of volumes that is seen with unsplit advection schemes that use face velocities such as Rider and Kothe (1998).

In order to back-trace the cell vertices, the velocities at these positions must be known. However, velocities are defined at cell faces. Interpolation is used to approximate the cell vertices velocities using cell faces. This introduces a problem in that the continuity equation ensures that the velocities at cell faces are divergence free which

cannot be said of the cell vertices. Thus, interpolation of velocities to cell vertices does not guarantee that that velocities remain divergence free. A correction to the flux volume is needed to ensure the conservation of mass during fluxing using this scheme of the geometric transport method. Refer to Owkes & Desjardins journal article for complete details of the unsplit geometric transport algorithm.

CHAPTER 5

RESULTS: TEST CASES

Reconstruction of a Randomly Placed Circle

In this test case, the PLIC normal approximation schemes are compared by testing the reconstruction of a simple circle that is placed randomly within a unit sized square domain. No transportation of the interface occurs and the exact interface is known. Therefore the best interface normal vector can be calculated for every cell in the computational domain. For the volume-of-fluid method, the best PLIC interface normal is the one that can minimize the distance between the exact interface and piecewise linear interface segment. After calculation of the best normal, the piecewise line segment is reconstructed such that it satisfies the condition that the area, or volume, under the piecewise linear segment is equal to the volume fraction of the cell.

In two-dimensions, the error in the normal approximation is calculated as the angle difference between the best normal and the normal obtained from the various normal approximation schemes. An ensemble average of L_1 error for 100 randomly placed circles of diameter 0.4 are calculated to remove any bias based on mesh positioning. In order to observe the convergence of each normal approximation scheme, the average L_1 error is obtained for various ratios of circle diameter, D , to mesh spacing, h . It is expected that as the circle interface is more accurately resolved using smaller cell volumes, the normal approximation error will decrease.

The ensemble average of the L_1 error norms are reported for each diameter, mesh spacing ratio on the next page.

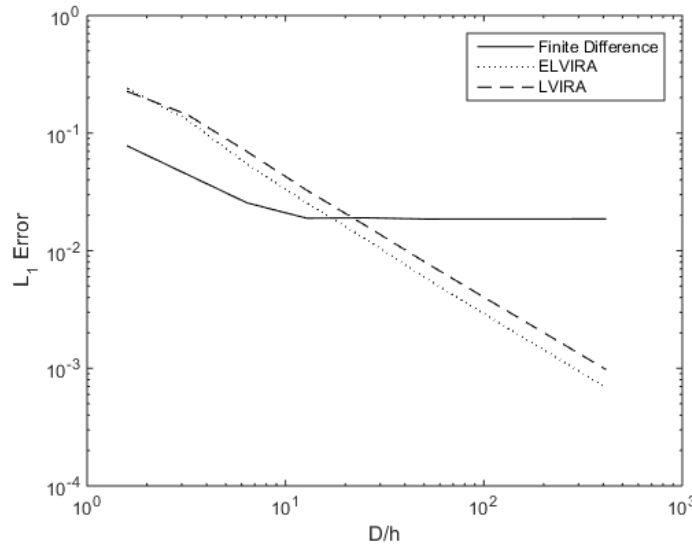


Figure 9 – Convergence of error for approximating the best normal when using the finite difference, ELVIRA, and LVIRA normal approximations.

D/h	Finite Difference	Convergence Rate	ELVIRA	Convergence Rate	LVIRA	Convergence Rate
1.6	7.75E-02		2.39E-01		2.26E-01	
3.2	4.45E-02	0.80	1.30E-01	0.88	1.44E-01	0.65
6.4	2.55E-02	0.80	5.44E-02	1.25	6.96E-02	1.05
12.8	1.89E-02	0.43	2.55E-02	1.09	3.27E-02	1.09
25.6	1.90E-02	-0.01	1.25E-02	1.04	1.62E-02	1.01
51.2	1.86E-02	0.04	5.81E-03	1.10	7.87E-03	1.04
102.4	1.86E-02	0.00	2.84E-03	1.03	3.93E-03	1.00
204.8	1.86E-02	0.00	1.40E-03	1.02	1.97E-03	1.00
409.6	1.86E-02	0.00	6.92E-04	1.02	9.82E-04	1.00

Table 1 – L_1 error norms of the normal approximation error when using the best normal described at the beginning of this section.

When approximating the interface normal of the randomly placed circles, the ELVIRA and LVIRA normal approximations schemes converge 1st order. The ELVIRA normal approximation scheme shows the lowest error to the best normal that was used for comparison calculations in this work. This is counter intuitive, since the LVIRA method uses an iterative process to minimize the error metric given by equation 6. This may be

an indication that equation 6 may not be the best metric for determining LVIRA, and ELVIRA.

The finite difference normal however performs below 1st order. It is also apparent that above a diameter to mesh spacing ratio of 12.8 the finite difference no longer converges. The reason for this is that as the circle is more accurately resolved, the interface curvature length scale becomes much, much larger than the cell size. As the curvature of the circle when viewed within a given cell volume approaches infinity, the normal can be approximated as a straight line. Pilliod & Puckett (2004), explain that the finite difference normal cannot accurately reconstruct a straight line when the interface passes through adjacent cell corners. Therefore, even when the mesh is infinitely fine, the finite difference scheme still cannot predict the exact normal interface. The ELVIRA and LVIRA formulations can however.

Zalesak's Disk Test Case Definition

The Zalesak's disk test case has been a standard test case that is used in evaluating interface capturing techniques like the Level-set, Marker-and-Cell, and Volume-of-Fluid methods. In these types of test cases shape errors and volume conservation errors are used as metrics to measure how well a method performs. The shape error is defined as.

$$E_{shape} = \sum_{p=1}^{N_{cv}} |C_{exact,p} - C_p|$$

It has been well documented in literature that the Volume-of-Fluid method has the ability to conserve mass to machine precision. For the consideration of space, volume conservation errors are not reported.

The Zalesak's Disk test case is the following: a notched disk is placed in a uniform sized box at a position of 0.5 units in the x-direction, and 0.75 units in the y-direction. The geometry of the disk is a 0.15 unit radius circle, with a notch that is 0.05 units wide, and 0.20 units high. Figure 10 below shows the Zalesak disk dimensions.




Figure 10 – Dimensions of the Zalesak notched disk that is used for this test case.

The velocity field in the unit-sized box is a prescribed rotational field that rotates the notched disk about the center of the box at an angular velocity of 1 radian per second.

$$u = 0.5 - y$$

$$v = x - 0.5$$

After 2π time units, the disk will have made one full rotation. The exact interface after one full rotation is exactly the initial condition. Thus, the shape errors and volume errors can be extracted using the initial values of the simulation.

For the Zalesak's disk test case, each scheme for advection and normal approximation was used. The final interface geometries are displayed on the following pages.

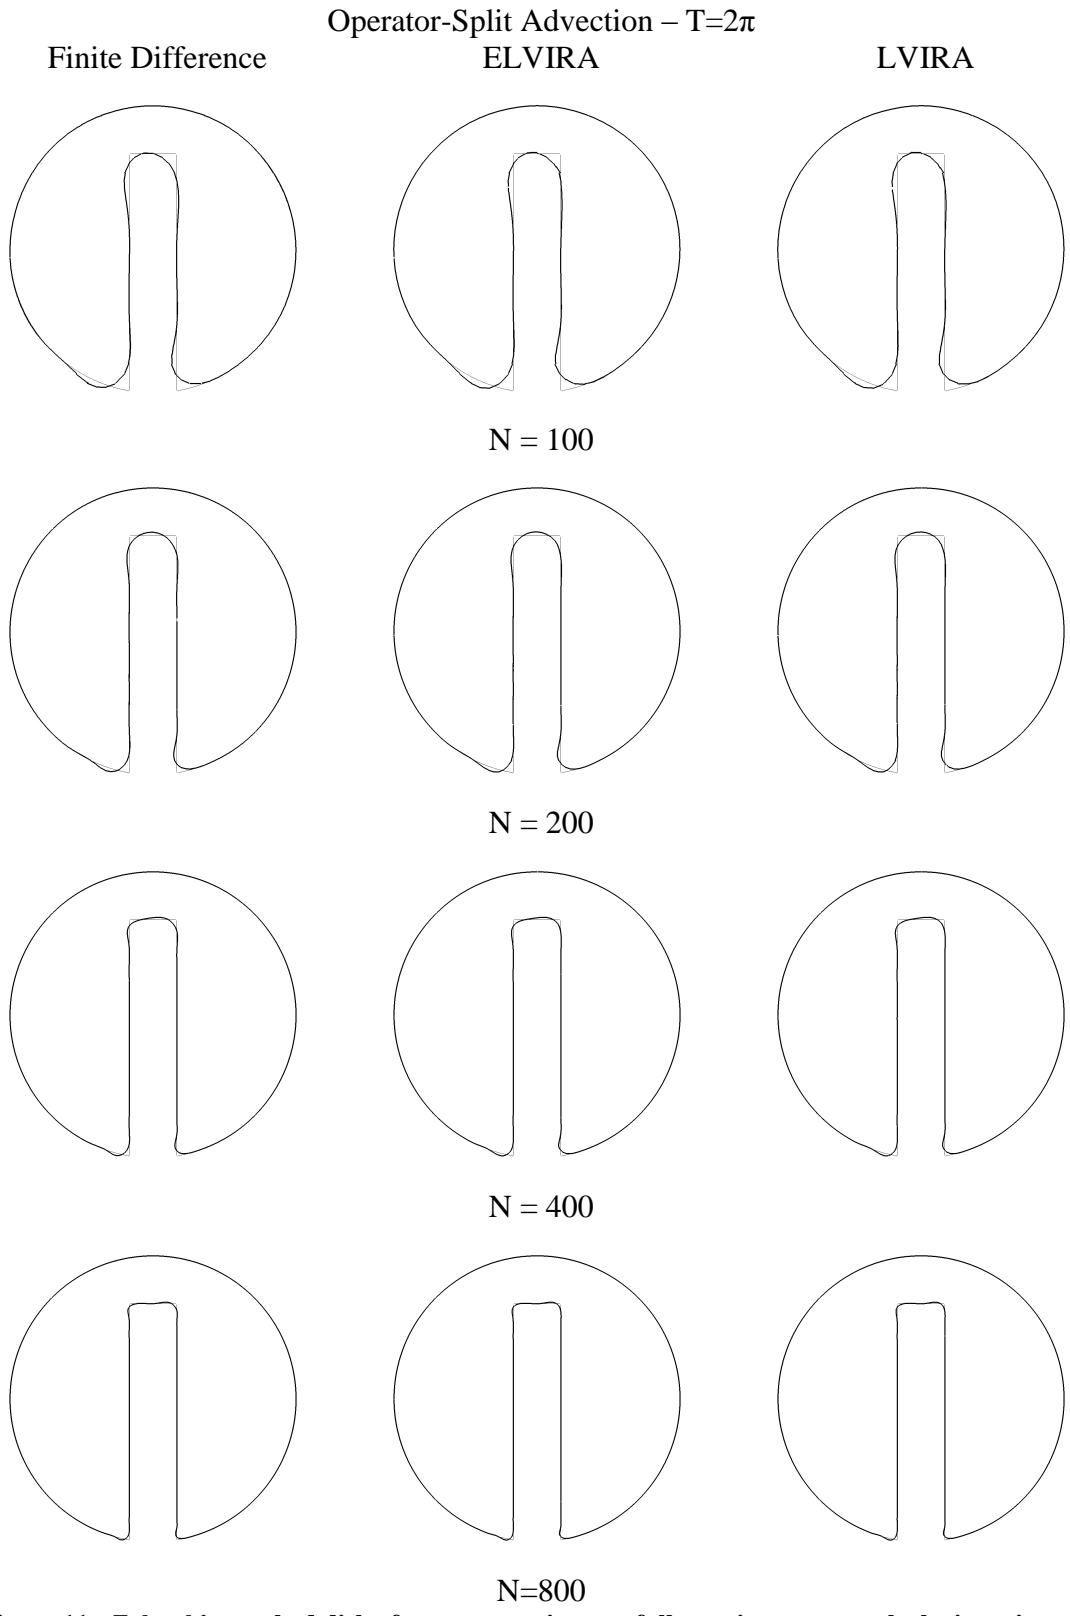


Figure 11 - Zalesak's notched disk after transporting one full rotation counter-clockwise using an operator-split non-rotational advection scheme.

Zalesak's Disk Results: Operator-Split Non-Rotational Advection

For the case of the Zalesak's disk the velocity vectors do not vary along their principle axes, that is $\partial u/\partial x = \partial v/\partial y = 0$. Since face velocities are equal on opposite faces of the cell, there is no rotation of PLIC surfaces. This leads to exactly the same fluxing operation for both of the operator-split advection methods tested in this work. For the sake of space, the depictions of the final shape figures will not be repeated for the second operator splitting advection technique.

In the Zalesak's disk test case using operator-split advection schemes there is graphically no difference between the final solutions when using different normal approximations. When observed very closely at the sharp edges of the notch, one can see very slight differences; these are shown in figure 12 for the upper right notched corner.




Figure 12 – Close-up view of the top right notch of Zalesak's disk. ELVIRA and LVIRA (grey) lie on the same curve, while the finite difference normal scheme is offset (black).

The shape errors for the Zalesak's disk test case also show very little differences between the different normal approximations. Table 2, shows the shape errors for the operator split, translational only advection scheme. Again, the shape errors for the second operator split advection scheme are omitted since any differences are miniscule and are only a result of fluctuations in truncation errors.




Figure 13 - Convergence of final result shape error for Zalesak's Disk test case.

Advection Scheme: Operator-Split Non-Rotational						
N	Finite Difference	Rate	ELVIRA	Rate	LVIRA	Rate
100	1.52E-03		1.44E-03		1.48E-03	
200	6.98E-04	1.13	6.23E-04	1.21	6.32E-04	1.23
400	3.23E-04	1.11	2.68E-04	1.22	2.75E-04	1.20
800	1.33E-04	1.28	1.01E-04	1.41	1.04E-04	1.41

Table 2 - Shape Errors for the operator-split non-rotational advection scheme.

The shape errors all show approximately first order, but at fine grid spacing this starts to improve as the rates become closer to 1.3 for the finite difference normal formulation, and closer to 1.4 for the least-squares normal approximations. Interestingly, the shape errors for the ELVIRA formulation are slightly lower than that of the LVIRA formulation observed for each grid spacing interval.

Next, the average computational time per time-step is shown for piece-wise linear interface reconstruction portion of the volume-of-fluid algorithm. Figure 14 depicts the

comparison between the finite difference normal, the ELVIRA normal, and LVIRA normal.

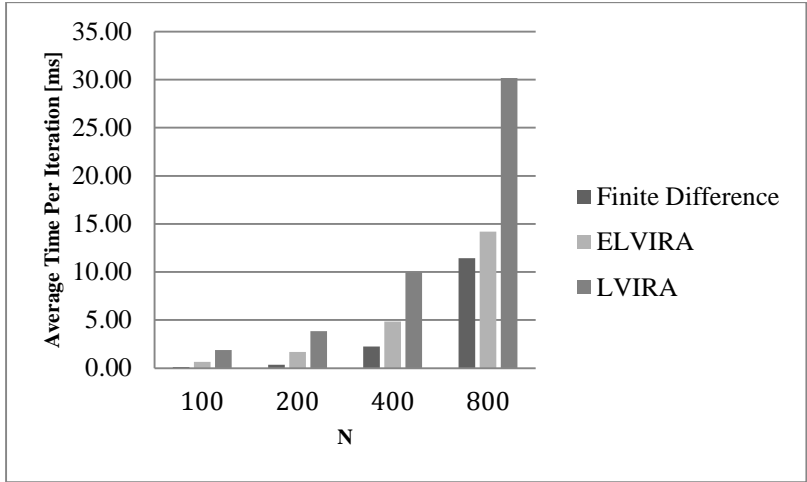


Figure 14 - Timing comparison between the Finite Difference, ELVIRA, and LVIRA normal approximation schemes.

Operator Split Non-Rotational							
		Finite Difference		ELVIRA		LVIRA	
N	N _{PLIC}	Timing [ms]	N _{PLIC}	Timing [ms]	N _{PLIC}	Timing [ms]	
100	1352	0.11	1352	0.66	1352	1.88	
200	3198	0.36	3199	1.68	3199	3.86	
400	6685	2.26	6686	4.84	6687	9.96	
800	13446	11.43	13448	14.20	13448	30.15	

Table 3 - Average time per iteration spent during reconstruction of the PLIC interface using the unsplit advection scheme for the Zalesak disk test case. N_{PLIC} represents the amount of cells that are reconstructed in the band structure.

The LVIRA normal approximation shows the longest time to approximate the interface normal directions. A large portion of this is due to the choice of using the ELVIRA normal as the initial guess. From this option, the LVIRA method starts out with the already costly ELVIRA method, and iterates to find the best normal. Timing information shows that the LVIRA method takes twice as long as the ELVIRA normal. In the test case of the Zalesak disk where interfaces do not become thin, the final result is largely indifferent from the choice of a normal approximation.

Zalesak's Disk Results: Unsplit Advection

The unsplit advection test cases are displayed on the next page. As compared to the operator-split advection methods, the unsplit advection scheme results from the Zalesak's disk test case show little differences in shape error. The shape errors for the unsplit test case are shown below in Table 4.




Figure 15 - Convergence of final result shape error for Zalesak's Disk test case using the second order advection scheme.




Advection Scheme: Unsplit						
N	Finite Difference	Rate	ELVIRA	Rate	LVIRA	Rate
100	1.45E-03		1.40E-03		1.44E-03	
200	6.72E-04	1.11	6.07E-04	1.20	6.20E-04	1.21
400	3.05E-04	1.14	2.56E-04	1.25	2.64E-04	1.23
800	1.25E-04	1.28	9.68E-05	1.40	1.00E-04	1.40

Table 4 - Shape Errors for the unsplit advection scheme applied to the Zalesak's Disk.




Finite Difference

Un-split Advection
ELVIRA




LVIRA






N = 100



N = 200



N = 400



N=800

Figure 16 - Zalesak's notched disk after transporting one full rotation counter-clockwise using an unsplit advection scheme.

The comparison of the shape errors in table 4 show very little difference when Zalesak’s disk is transported using the unsplit advection scheme. The ELVIRA and LVIRA normal approximations show a slight advantage over the finite difference normal formulation. This is also apparent in the operator-split advection schemes as well.

The average computational time spent during the piecewise linear interface reconstruction for the Zalesak’s disk test case with the unsplit advection scheme is visually depicted below in figure 17. Table 5, shows the timing values given in milliseconds.




Figure 17 - Timing comparison of the Finite Difference, ELVIRA, and LVIRA normal schemes.

Advection Scheme: Unsplit						
N	N _{PLIC}	Finite Difference	N _{PLIC}	ELVIRA	N _{PLIC}	LVIRA
		Timing [ms]		Timing [ms]		Timing [ms]
100	1130	0.101	1130	0.554	1131	1.120
200	2599	0.303	2599	0.917	2600	1.985
400	5245	0.914	5245	2.064	5245	4.357
800	10543	4.414	10542	6.004	10542	12.723

Table 5 – Average time per iteration spent during reconstruction of the PLIC interface using the unsplit advection scheme for the Zalesak disk test case. N_{PLIC} represents the amount of cells that are reconstructed in the band structure.

In the computation of the Zalesak's disk test case using unsplit advection, the time spent during the piecewise linear interface reconstruction step is less than when using the operator-split methods. In general, the number of cells that need to be geometrically reconstructed was lower than in operator-split advection, which resulted in a reduction in interface reconstruction time per iteration. The reason for this is that only the cells in the transportation band need to be reconstructed. Since the unsplit method is not prone to float-sum like those created in an operator split advection scheme, there is no need to reconstruct interior cells in the fluid because there are no cells that need to be reconstructed.

The general trend between the different normal schemes is still apparent. The finite difference method shows the lowest reconstruction times. The ELVIRA scheme shows higher reconstruction timing than that of the finite difference scheme. It is approximately five times slower at a mesh sizing of 100x100 cells, and 1.5 times slower for a mesh sizing of 800x800. The LVIRA method shows the slowest reconstruction times which took 2 times as long as the ELVIRA scheme, and 3 times as long as the finite difference scheme for the mesh sizing of 800x800 cells.

Zalesak's Disk Results: Advection Timing Comparisons

Lastly, for the Zalesak's disk test case the average time per iteration of the advection step are of interest. The number of cells on which advection of volume-of-fluid volume fraction occurs is nearly constant regardless of the choice of normal approximation. Therefore, the advection timing displayed is the average timing over the three normal approximation schemes.

For timing, the average advection times are shown for each advection step, for each mesh size. These are depicted in figure 18 and shown in table 6.

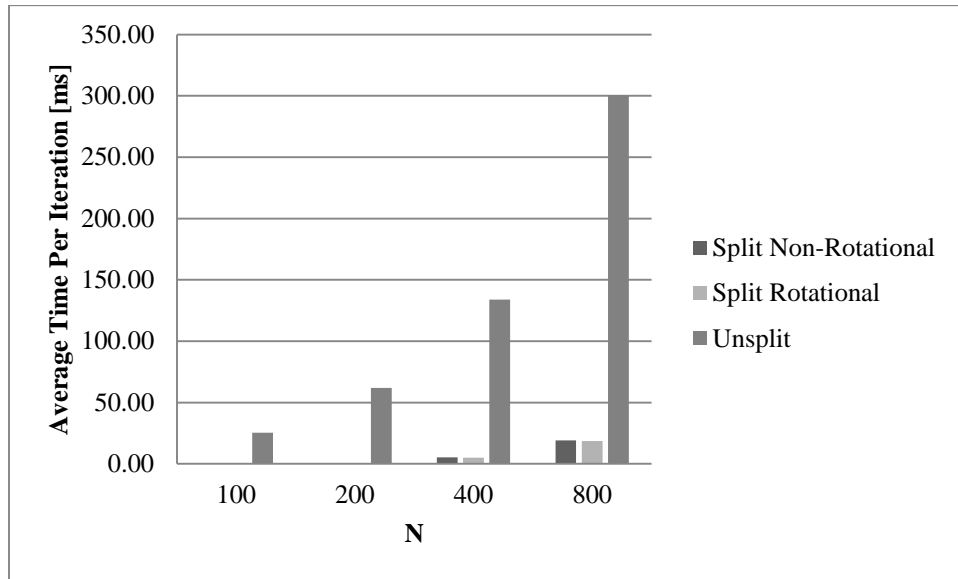


Figure 18 - Comparison of average time spent during the advection step of the volume-of-fluid algorithm for the Zalesak’s disk test case.

Advection Timing [ms]			
N	Split Non-Rotational	Split Rotational	Unsplit
100	0.13	0.15	25.42
200	0.53	0.56	61.75
400	5.19	5.05	133.76
800	19.12	18.63	300.49

Table 6 – Average computational time per iteration for each advection scheme tested using Zalesak’s Disk.

The most obvious observation is that the 2nd order conservative unsplit advection method is extremely costly compared to the operator-split methods. On the finest mesh size tested, the average time during the transportation step using the unsplit advection scheme was approximately 15 times higher as compared to the operator-split methods. The timing comparison also shows that the cost of both operator-split methods are almost exactly the same for every mesh size.

Circle in a Deformation Field Definition

The second test case that is used to measure the performance of the various normal and advection schemes is the two dimensional column that is placed inside of a deformation velocity field. For this test case a circle of radius 0.15 units is placed in a unit sized square domain with its center at 0.5 units in the x-direction and 0.75 units in the y-direction. A velocity field is prescribed according to the following stream function over the course of $T=8$ time units.

$$\Psi(x, y) = \frac{1}{\pi} \sin^2(\pi x) \cos^2(\pi y) \cos\left(\frac{\pi t}{T}\right)$$

The velocity field is periodic in time, therefore after $t > \frac{T}{2}$ the flow field is reversed. The purpose of this time dependent deformation field is to deform the circle into a stretched ligament. Figure 19 below shows the deformation of the circle interface over the course of 4 time units. After 8 time units, the initial condition should completely reform.




Figure 19 - Time progression of the circle placed in a deformation field test case. As time progresses, the interface follows from left to right until deformed to a stretched ligament. When the circle is completely deformed, the flow field is reversed, and follows right to left until it reforms the initial circle.

Circle in a Deformation Field Results: Non-Rotational Advection

The depictions of the fluid interface are shown on the next page for the deformation column test case at the half mark of the interface capturing simulation. Starting with the mesh size of 64^2 , the finite difference approximation clearly shows the worst performance. The ligament is broken into many segments and starts nearly 50 percent of the length of the stretched ligament. The ELVIRA normal approximation shows better stability in holding a continuous interface when the interface becomes thin. At the top of the circle ligament the interface is rough using the ELVIRA normal, and very small discontinuities in the PLIC reconstruction are visible. The LVIRA normal approximation scheme shows the best performance in maintaining the thin interface at the end of the ligament. In total there are two major ligament breakups in which the normal becomes rounded and the ligament becomes separated. There is a third break in the fluid interface, however close inspection shows the interface normal directions are still continuous. The reason there is a visual break across the interface is that interface falls below the thickness of the cell and only one PLIC interface is drawn per cell. If the ligament were to continue to stretch, this would be the next breaking point in the ligament since the ligament thickness in this region falls below one cell thickness, h .




In the 128^2 size mesh, the length of the ligament for each of the test cases of the normal approximation scheme becomes longer. The finite difference normal scheme, again, shows the most droplets at the tail of the ligament. It has five droplets that have broken off from the ligament. The ELVIRA normal continues to do better than the finite difference normal, and the roughness in the interface reconstruction smoothens.

Operator-Split Advection – Non-Rotational – $T = 4$




Finite Difference

ELVIRA




LVIRA






$N = 64$



$N = 128$



$N = 256$



$N=512$

Figure 20 - Resulting fluid interface after $t = 4$ when the circle is completely stretched into a ligament using the operator-split non-rotational advection scheme for transportation.

The LVIRA normal again shows the lowest amount of breakup, in which only one segment of the ligament, breaks off.

This trend continues as the mesh becomes finer in the 256^2 mesh. Almost all breakups disappear in the finite difference and ELVIRA schemes, and the LVIRA scheme shows no breakup at all. In the finest mesh tested, when 512^2 , all the test cases using the different normal approximations are visual similar, and no differences can be seen.

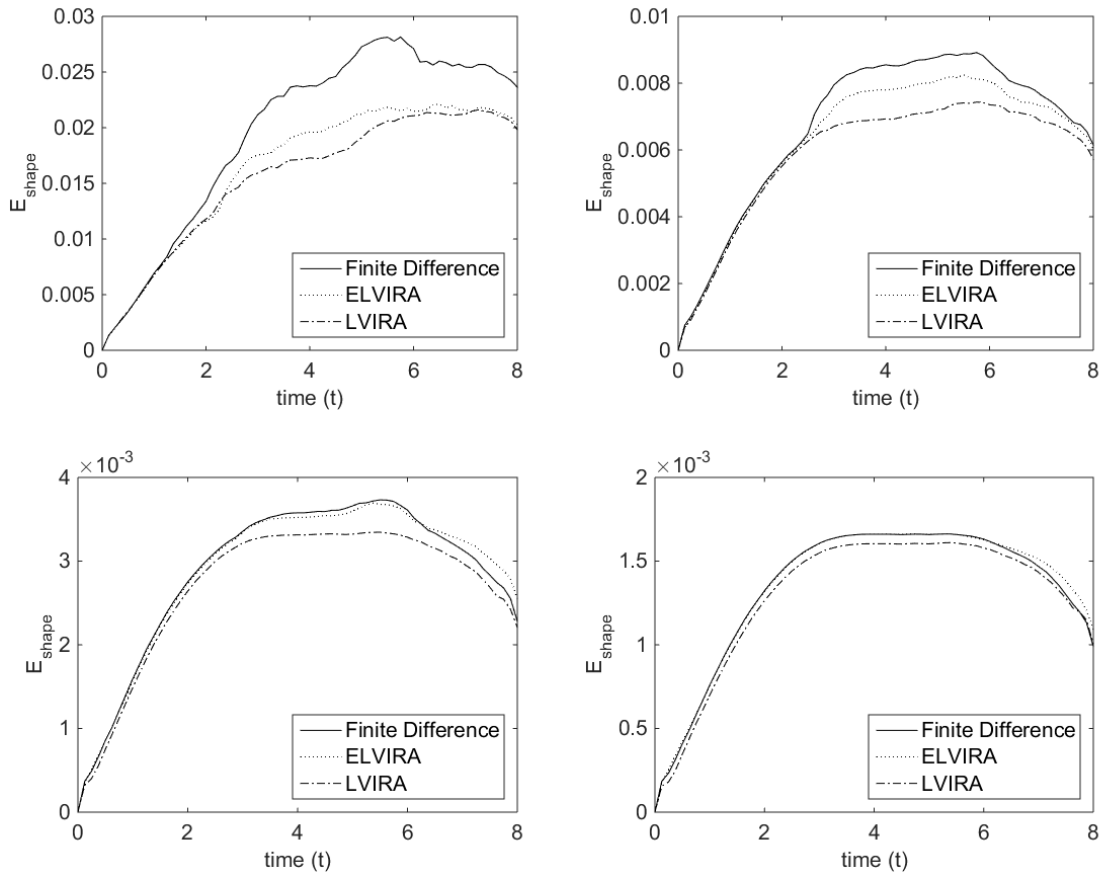


Figure 21 - Shape errors plotted over the entire simulation. (a) Top left: $N=64$, (b) top right: $N=128$, (c) bottom left: $N=256$, (d) bottom right: $N=512$.

When the shape error of the interface is plotted over the entire simulation, as is shown in figure 21, it is easier to see that the LVIRA normal approximation scheme provides the lowest shape errors through the first half of the simulation. As the

transportation reverses, there is apparent error cancellation upon reformation of the circle. Actual shape error numbers are not given since the plots give approximate shape errors based on the results of a high mesh resolution test case. The curves in figure 21, are only to give better insight to what is happening throughout the simulation.

Next, the fluid interface depictions for the final resulting solution are compared. The depictions of the fluid interface are shown on the next page. The most notable details in the final fluid interfaces is the overshoot that appears in each of the test cases. It appears as if the simulation had run for one or two extra time-steps, however these are the result of the advection splitting errors. Secondly, all of the methods show a buildup of float-sum within the interior of the fluid region. This is also a result of the advection splitting errors, where the splitting of the advection operation results in the fluxing of the volume fraction during separate non-divergence free advection steps. As a result, some cells retain a volume fraction either lower than 0, or higher than 1.




Starting at the top, $N=64$, the final results for each of the normal approximation schemes show high amounts of error. It is difficult to judge visually which method performs the best. When the mesh is refined, the ELVIRA and LVIRA start to show the more volume retained within the initial circle. At 256^2 , the LVIRA method nearly fully covers the entire initial condition, while the ELVIRA and finite difference methods still have slight lumps. Again at 512^2 , the differences are unapparent and the solution is nearly exact disregarding the errors due to float-sum.

Operator-Split Advection – Non-Rotational – $T = 8$




Finite Difference

ELVIRA




LVIRA






$N = 64$



$N = 128$



$N = 256$



$N=512$

Figure 22 - Final solution fluid interface when $t = 8$ after the circle is completely deformed and reformed using the operator-split non-rotational advection scheme for transportation.

The performances of the individual methods are more easily observed through the calculation of the final result shape error. The convergence of the shape errors is shown in figure 23 below. Following the convergence plot is table 7 containing the values of the final shape errors.




Figure 23 - Convergence history of the shape error in the deformation circle test case using the operator-split non-rotational advection scheme.

Advection Scheme: Operator-Split Non-Rotational						
N	Finite Difference	Convergence Rate	ELVIRA	Convergence Rate	LVIRA	Convergence Rate
64	2.362E-02	-	2.030E-02	-	1.982E-02	-
128	6.167E-03	1.94	6.038E-03	1.75	5.752E-03	1.78
256	2.329E-03	1.40	2.567E-03	1.23	2.252E-03	1.35
512	1.027E-03	1.18	1.106E-03	1.21	1.041E-03	1.11

Table 7 - Shape errors for the operator-split non-rotational advection scheme applied to the circle in a deformation field test case.

On the coarsest mesh tested, the LVIRA method showed the lowest shape errors. This remained the same for the 128^2 and 256^2 sized cell meshes. However for this comparison using the operator split non-rotational advection scheme, the finest mesh shape errors

showed that the finite difference had the lowest shape errors. The LVIRA method showed the second lowest shape errors on the finest cell mesh, and finally the ELVIRA method showed the highest errors. However, the shape errors between the comparisons are largely the same on the finest mesh, and the stacking of the shape error comparisons could be due to the amount of float-sum in the models. The ELVIRA method showed the highest shape error, but the final result using ELVIRA also had the highest amount of float-sum.




Figure 24 - Timing comparison of the Finite Difference, ELVIRA, and LVIRA normal schemes using the operator-split non-rotational advection scheme for the circle in a deformation field test case.

Advection Scheme: Operator-Split Non-Rotational						
N	Finite Difference		ELVIRA		LVIRA	
	N_{PLIC}	Timing [ms]	N_{PLIC}	Timing [ms]	N_{PLIC}	Timing [ms]
64	1784	0.147	1776	0.974	1775	2.610
128	4344	0.629	4340	2.521	4329	6.489
256	10742	4.040	10829	8.834	10689	22.038
512	27274	15.425	28127	25.539	27094	56.658

Table 8 - Average time per iteration spent during reconstruction of the PLIC interface using the operator-split non-rotational advection scheme. N_{PLIC} represents the amount of cells that are reconstructed in the band structure.

The average timing for the PLIC operations is compared next. Figure 24 shows a bar graph comparison for the timing between the individual normal approximation schemes. Table 8 shows the expected trend that finite difference normal scheme requires the lowest computational time. The ELVIRA follows, generally taking twice as long to reconstruct the interface. Finally the LVIRA method clearly takes the longest to reconstruct the fluid interface. If one were to compare the shape error per computational time, the finite difference method would clearly win, since the shape errors are generally of the same order, while the computational time difference is much larger.

Circle in a Deformation Field Results: Rotational Advection




The operator-split rotational advection scheme is tested next. On the next page are the visuals of the fluid interface at the half mark of the simulation when the circle is completely deformed into the ligament. The deformed ligament is visually very similar to that of when using the non-rotational advection scheme. Throughout the study for this advection scheme, the finite difference and ELVIRA normal approximation schemes are very similar to the non-rotational advection counter-parts. However, the results from the LVIRA method show some differences. Firstly, in the 64^2 mesh, the tail of the circle ligament breaks into four small droplets. The length of the ligament when using LVIRA is still clearly the longest. In the 128^2 mesh, the ligament also shows no breakup at the tail when using LVIRA. This is slightly better than the operator-split non-rotational advection scheme where using LVIRA had one small droplet pinch off. This is also much better than the two other normal approximations shown on the next page for mesh sizes 128^2 for finite difference and ELVIRA. At the last stage, the mesh is fine enough that there are no visible differences.

Operator-Split Advection – Rotational – $T = 4$




Finite Difference

ELVIRA




LVIRA






$N = 64$



$N = 128$



$N = 256$



$N=512$

Figure 25 - Resulting fluid interface after $t = 4$ when the circle is completely stretched into a ligament using the operator-split rotational advection scheme for transportation.

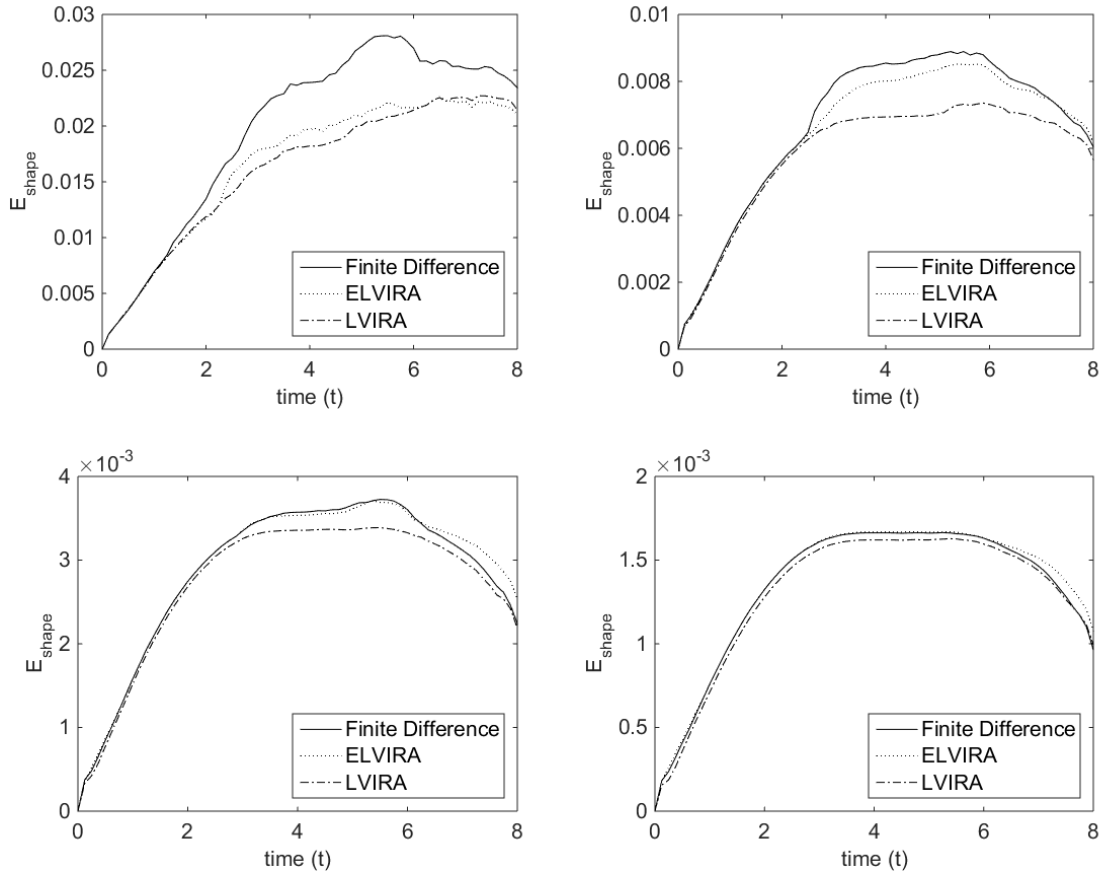


Figure 26 - Approximate shape errors plotted over the entire simulation. (a) Top left: $N=64$, (b) top right: $N=128$, (c) bottom left: $N=256$, (d) bottom right: $N=512$.

Similarly to the operator-split non-rotational method study, the approximate shape errors can be plotted over the entire simulation. The curves in figure 26 show similarity to the operator-split non-rotational shape error curves in that the LVIRA scheme achieves lower shape error through the majority of the simulation. Toward the end of the simulation, upon reformation of the circle, error cancellation occurs. In the 64^2 test case, the ELVIRA normal actually dips below the LVIRA method in the final seconds of the simulation. One can approximate when the ligaments experience break up by when there is a sudden jump in shape error. This is most apparent in the 128^2 mesh shortly after 2 time units.

The visuals for the final results from the deformation circle using the three normal comparisons are shown on the next page. The first observation is the amount of float-sum that is apparent in each of the final results. The ELVIRA method shows the highest amount of float-sum, compared to the results using the finite difference and LVIRA approximations. Still, the amount of float-sum that is apparent in the results is unsatisfactory. This is because the amount of float-sum increases the number of cells that the algorithm perceives that must be geometrically reconstructed. In doing so, this increases the number of cells that are in the transportation band. Overall it increases the number of cells on which the volume-of-fluid algorithm is used for transportation of the interface, leading to an increase in computational cost. Secondly, the overshoot is still apparent due to the operator splitting errors.




Observing figure 27, and starting with the 64^2 mesh the finite difference method has a large amount of area that is outside of the original circle. Between the ELVIRA and LVIRA, it is harder to tell what has the better performance. Moving to 128^2 , the circle is better reformed going from left to right along finite difference, ELVIRA, and LVIRA. Continuing to 256^2 , the LVIRA method is nearly the initial condition, while finite difference and ELVIRA still have traces of lumps on the upper edge of the circle. Finally, at 512^2 , the differences in the final formation are indistinguishable.

Operator-Split Advection – Rotational – $T = 8$




Finite Difference

ELVIRA




LVIRA






$N = 64$



$N = 128$



$N = 256$



$N=512$

Figure 27 - Final solution fluid interface when $t = 8$ after the circle is completely deformed and reformed using the operator-split rotational advection scheme for transportation.

Final shape errors for the rotational operator split advection scheme are displayed next. Figure 28 shows the convergence plot for every mesh tested. Table 9, contains the information for shape errors at the final time-step.




Figure 28 - Convergence history of the shape error in the deformation circle test case using the operator-split rotational advection scheme.

Advection Scheme: Operator-Split Rotational						
N	Finite Difference	Convergence Rate	ELVIRA	Convergence Rate	LVIRA	Convergence Rate
64	2.349E-02	-	2.107E-02	-	2.150E-02	-
128	6.125E-03	1.94	6.233E-03	1.76	5.725E-03	1.91
256	2.304E-03	1.41	2.566E-03	1.28	2.252E-03	1.35
512	1.012E-03	1.19	1.105E-03	1.22	1.044E-03	1.11

Table 9 - Shape errors for the operator-split non-rotational advection scheme applied to the circle in a deformation field test case.

In figure 28, there is hardly any difference in the convergence rates using the different normal schemes. Initially every method converges close to second order. As the mesh becomes finer the convergence rates slow to near first order. The LVIRA method usually shows the lowest shape error values until the finest mesh tested where the finite

difference normal achieves the lowest shape error. Referring back to figure 26, this is due to the error cancelation that occurs near the end of the simulation. For a majority of the simulations, the finite difference method shows the highest shape errors.




Figure 29 - Timing comparison of the Finite Difference, ELVIRA, and LVIRA normal schemes using the operator-split rotational advection scheme for the circle in a deformation field test case.

Advection Scheme: Operator-Split Rotational						
N	Finite Difference		ELVIRA		LVIRA	
	N_{PLIC}	Timing [ms]	N_{PLIC}	Timing [ms]	N_{PLIC}	Timing [ms]
64	2038	0.150	2027	1.098	2011	3.059
128	6380	0.617	6339	2.885	6319	7.078
256	20987	5.591	20958	9.849	20616	25.195
512	67931	18.722	68538	29.397	67721	64.847

Table 10 - Average time per iteration spent during reconstruction of the PLIC interface using the operator-split rotational advection scheme. N_{PLIC} represents the amount of cells that are reconstructed in the band structure.

Timing information for the operator-split rotational advection scheme is shown above. Observing table 10, the finite difference scheme performs the fastest. The ELVIRA method is second fastest, and the LVIRA method shows the highest cost. The order of magnitude in the number of float-sun can be perceived from table 10 in the number of PLIC elements column. The ELVIRA method is confirmed to have the most

amount of float-sum since the N_{PLIC} elements indicate the number of elements that are geometrically reconstructed.

Circle in a Deformation Field Results: Unsplit Advection

The unsplit advection scheme is the last advection scheme tested for this test case. Figure 30 on the next page shows the fluid interfaces at the half mark of the test case simulation. Overall when compared to both of the operator-split advection schemes, many of the features, such as interface shape, remain prevalent.

On the 64^2 mesh, the ligament tail breakup is largely unchanged, despite using the unsplit advection scheme. The finite difference normal scheme still exhibits large amounts of artificial fluid breakup that is present in half of the length of the ligament. This is due to the sensitivity of the finite difference normal scheme to the thickness of the interface. When the interface thickness becomes small enough that there are two interfaces present in the 3×3 cell stencil, the finite difference normal approximation becomes inaccurate. This directly affects the fluxing operation. In the ELVIRA normal approximation, the ligament is preserved longer, much like in the test cases using the operator split advection schemes. The LVIRA normal scheme still shows the highest performance when it comes to preserving the thin interface.

In the 128^2 mesh the ligaments are longer, and LVIRA shows no signs of breakup. ELVIRA shows breakup, and the finite difference scheme shows the most breakup of all the normal schemes.

Continuing to 256^2 , most of the ligament breakup has disappeared, however there are still two small pinch offs at the tip of the tail. At the final 512^2 , the mesh resolution is fine enough that each method performs nearly the same.

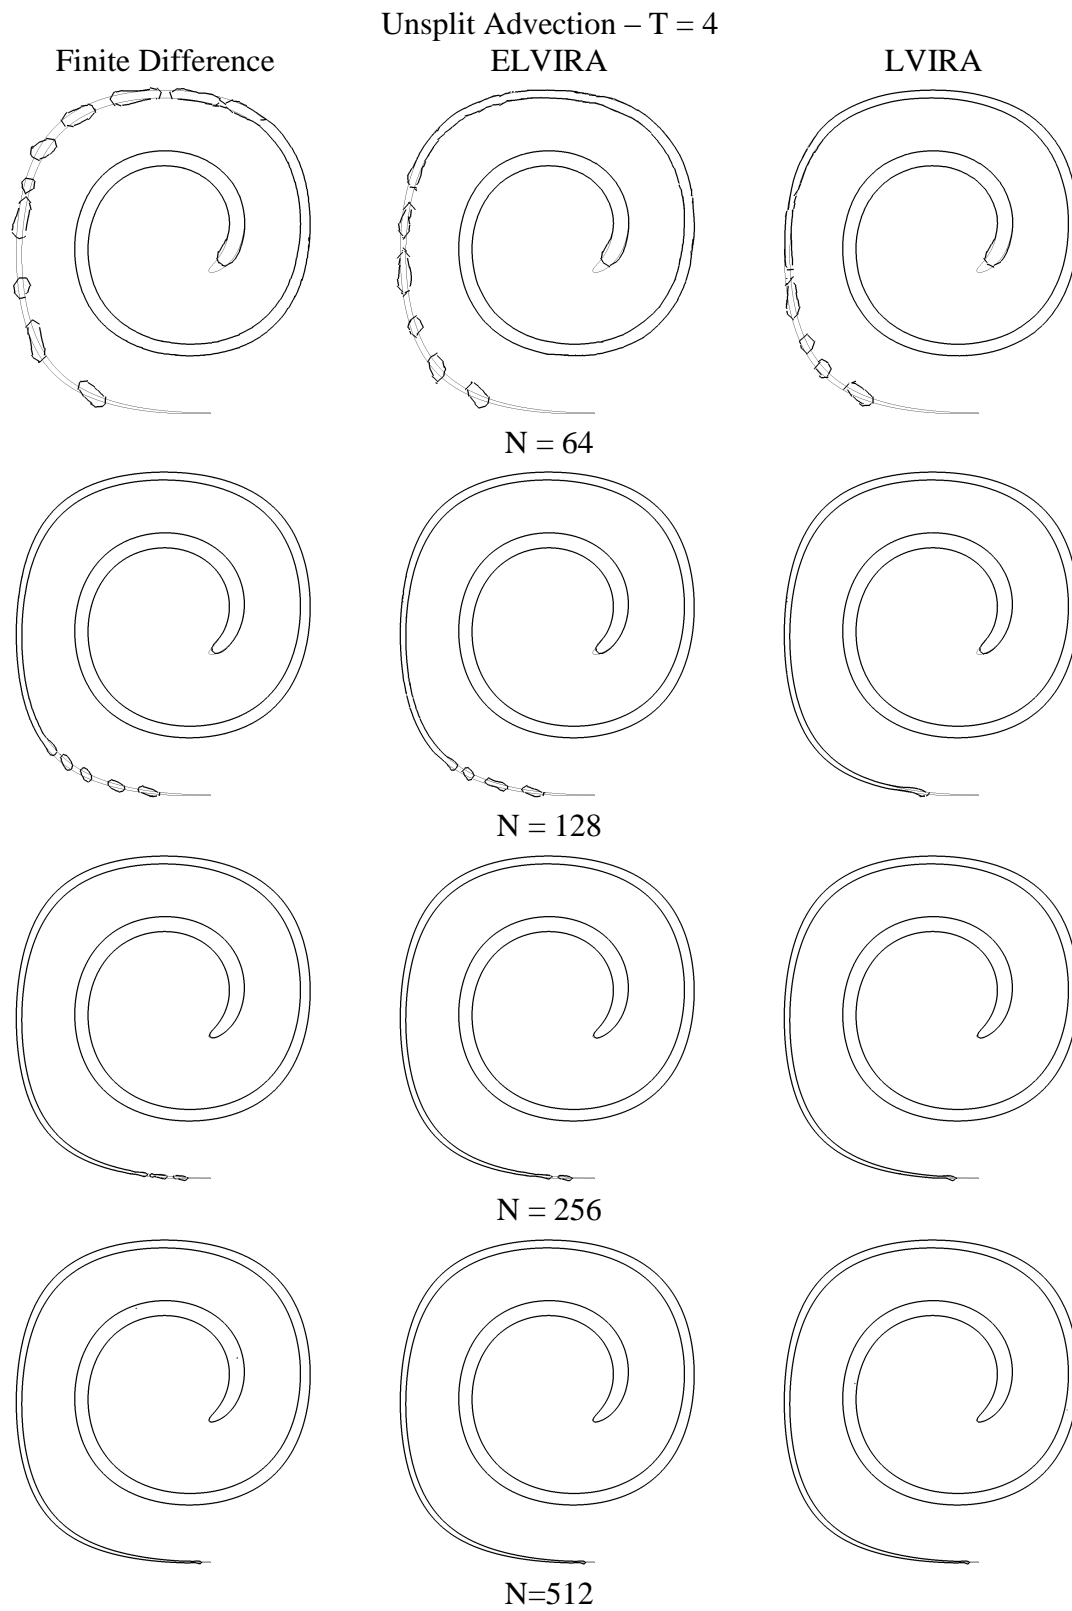


Figure 30 - Resulting fluid interface after $t = 4$ when the circle is completely stretched into a ligament using the unsplit advection scheme for transportation.

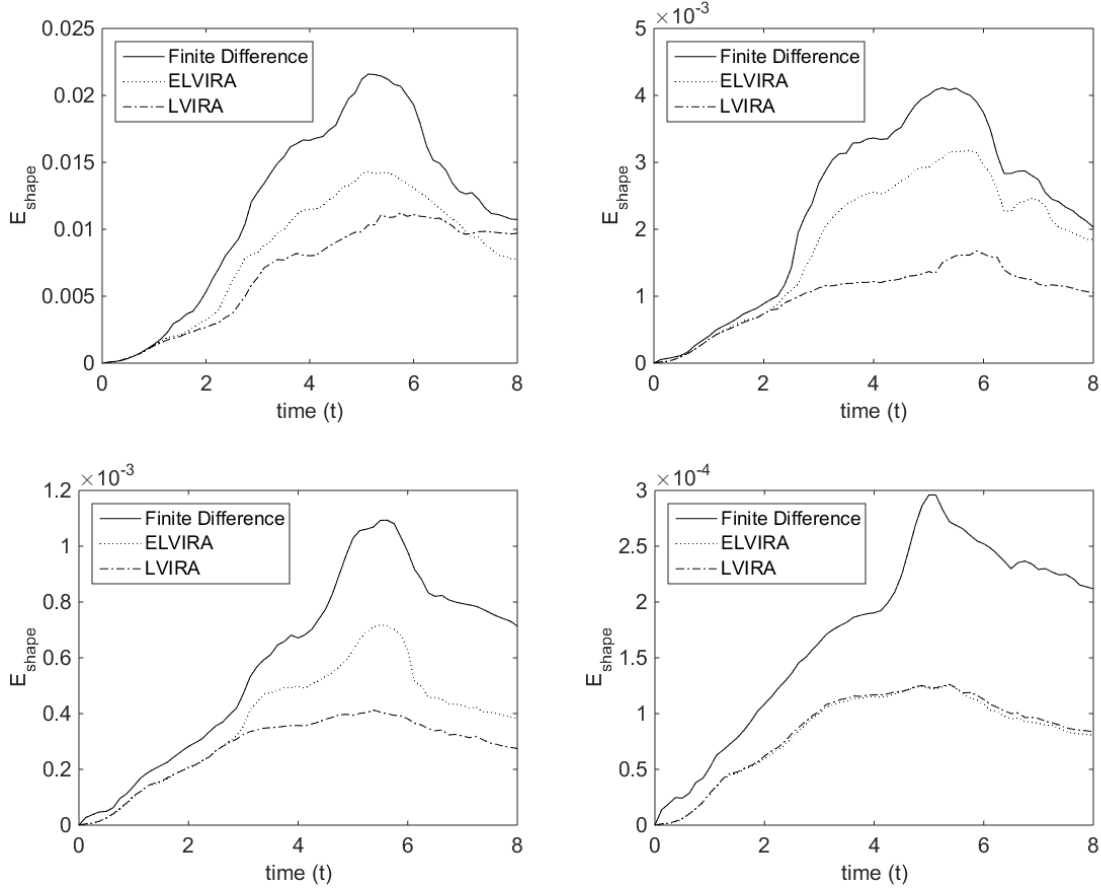


Figure 31 - Approximate shape errors plotted over the entire simulation. (a) Top left: $N=64$, (b) top right: $N=128$, (c) bottom left: $N=256$, (d) bottom right: $N=512$.




The shape errors throughout the simulation, figure 31, show much more spread between the different normal approximation schemes. It can be seen that the finite difference shows much higher shape errors over the simulation. In figure 31a, it is clearly seen that the error cancellation properties of the time reversal field cause the finite difference and ELVIRA normal approximation shape errors to have significant error reduction when the ligament reforms into a circle. The latter results in a final shape error that is lower than seen using LVIRA.

The fluid interfaces at the final time-step are shown on the next page in figure 32 for the unsplit advection scheme in the deformation circle test case. The first most obvious observation is the absence of float-sum in the final solution.




Finite Difference

Unsplit Advection – T = 8
ELVIRA




LVIRA






N = 64



N = 128



N = 256



N=512

Figure 32 - Final solution fluid interface when $t = 8$ after the circle is completely deformed and reformed using the unsplit advection scheme for transportation.

The removal of float-sum creates fluid interfaces that are much cleaner visually and the result looks much cleaner. The absence of float-sum also will provide more accuracy when coupled to a Navier-Stokes solver. Any types of float-sum will cause problems and inaccuracies for the calculation of surface tension. The float-sum will begin to cause artificial bubbles within the fluid geometry.

The next observation is the absence of the overshoot that was apparent in the operator split advection schemes. At the lowest mesh resolution, there is a large perimeter of the circle that lies on the initial condition. The removal of the operator-splitting step removed the splitting errors that cause the overshoot.

Comparing the results from the different normal schemes, and starting at $N=64$, the finite difference method shows the largest area outside of the initial condition circle. The ELVIRA and LVIRA methods are similar, but the ELVIRA method has the smaller area on the outside of the circle. Continuing to $N=128$, it can be seen when moving from left to right, that the LVIRA method performs the best. The ELVIRA method is second, and the finite difference scheme is last. This trend continues on into the $N=256$ mesh. By this mesh refinement level, the LVIRA method shows nearly no error. The ELVIRA method has small lumps in the upper portion of the circle. The finite difference method shows large bumps in the same region as the ELVIRA method.




Figure 33 - Convergence history of the shape error in the deformation circle test case using the unsplit advection scheme.

Advection Scheme: Unsplit						
N	Finite Difference	Convergence Rate	ELVIRA	Convergence Rate	LVIRA	Convergence Rate
64	1.073E-02	-	7.772E-03	-	9.700E-03	-
128	2.031E-03	2.40	1.850E-03	2.07	1.053E-03	3.20
256	7.136E-04	1.51	3.835E-04	2.27	2.746E-04	1.94
512	2.118E-04	1.75	8.063E-05	2.25	8.377E-05	1.71

Table 11 - Shape errors for the unsplit advection scheme applied to the circle in a deformation field test case.

On the coarsest mesh tested the ELVIRA method performs the best. LVIRA has the second highest error, and finite difference performs the worst, with the highest error. On the 128^2 and 256^2 meshes, LVIRA performs the best. At the 512^2 mesh size, the ELVIRA method performs the best. It can be seen that the convergence rate when using unsplit advection was near second order as compared to the operator-split advection schemes that showed near first order convergence. Secondly, the unsplit advection scheme showed much more deviance using the different normal schemes. Interestingly, the ELVIRA

method maintains above 2nd order convergence. It is likely for higher mesh resolutions than 512², the shape errors for ELVIRA and LVIRA will converge to the same solution. This is because as the relative curvature of the interface approaches infinity, ELVIRA and LVIRA approach the same approximation. Ideally, when the curvature reaches infinity, the interface becomes a straight line locally for each 3x3 cell stencil where the normal is approximated. For straight lines both ELVIRA and LVIRA produce the exact normal (Pilliod & Puckett, 2004).

The PLIC reconstruction timing is next to be analyzed. Figure 34 contains a graph on how the average reconstruction timing compares between the normal comparisons.

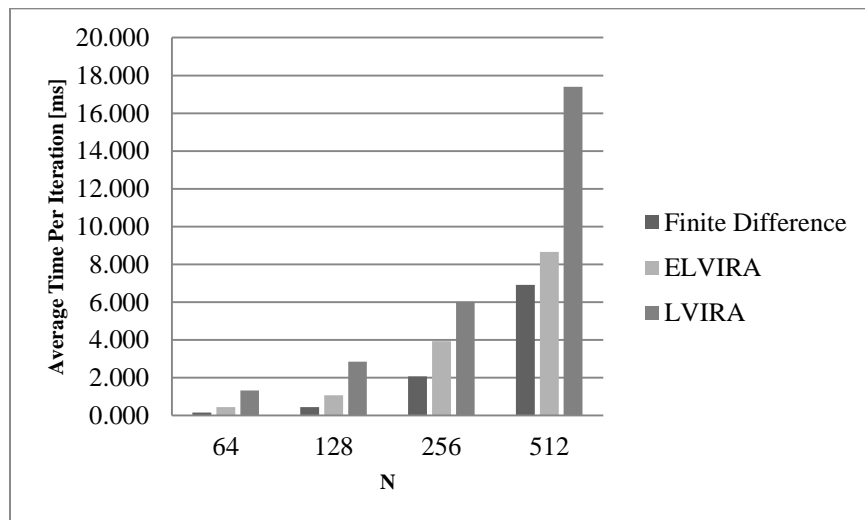


Figure 34 - Timing comparison of the Finite Difference, ELVIRA, and LVIRA normal schemes using the unsplit advection scheme for the circle in a deformation field test case.

Advection Scheme: Unsplit						
N	Finite Difference		ELVIRA		LVIRA	
	N _{PLIC}	Timing [ms]	N _{PLIC}	Timing [ms]	N _{PLIC}	Timing [ms]
64	1520	0.162	1512	0.447	1511	1.325
128	3577	0.441	3576	1.068	3567	2.853
256	8644	2.070	8641	3.936	8634	6.013
512	19257	6.913	19206	8.658	19224	17.406

Table 12 - Average time per iteration spent during reconstruction of the PLIC interface using the unsplit advection scheme. N_{PLIC} represents the amount of cells that are reconstructed in the band structure.

The comparison in figure 34 shows the expected trend. The LVIRA method is the most costly, taking anywhere from 8 times as long as the finite difference in the coarsest mesh and 2.5 times as long seen in the finest mesh used. The ELVIRA method is the second most costly. It takes 2.75 times as long as the finite difference method on the coarsest mesh and 1.25 times as long as the finite difference method on the finest mesh. When the shape errors are taken into consideration, the shape error when using the finite difference method at the finest mesh is 2.6 times larger than the ELVIRA method. So for the ELVIRA method the computational cost is worth the reduction in error.

Circle in a Deformation Field Results: Advection Timing Comparison




Figure 35 - Comparison of average time spent during the advection step of the volume-of-fluid algorithm for the circle in a deformation field test case.

N	Advection Timing [ms]		
	Split Non-Rotational	Split Rotational	Unsplit
64	0.199	0.214	27.773
128	1.225	1.445	63.509
256	8.087	8.623	164.620
512	29.199	32.786	423.080

Table 13 – Average computational time per iteration for each advection scheme tested in the circle in a deformation field test case.

For the circle in a deformation field test case it is useful to know how the advection schemes compare. This comparison is given in figure 35. The unsplit advection scheme by far takes the longest to complete. It takes anywhere from 100 times as long on the coarsest mesh to 20 times as long on the finest mesh. The two operator split advection schemes took nearly the same time to compute the advection step.

Sphere in a Deformation Field Test Case Definition

The last test case used in this work is the sphere placed in a deformation field test case. The test case is three-dimensional and is the following. A sphere of radius 0.15 units is placed inside of a unit-sized box. The center of the sphere is placed at coordinates, $x = 0.35$, $y = 0.35$, and $z = 0.35$. The deformation field prescribed is the following,

$$u = 2 \sin^2(\pi x) \sin(2\pi y) \sin(2\pi z) \cos(\pi t/3)$$

$$v = -\sin(2\pi x) \sin^2(\pi y) \sin(2\pi z) \cos(\pi t/3)$$

$$w = -\sin(2\pi x) \sin(2\pi y) \sin^2(\pi z) \cos(\pi t/3)$$

The test case is similar to the circle in a deformation field in that the deformation field is periodic. Halfway through the simulation, the flow field is reversed and the sphere will reform into the initial condition. At the end of the simulation shape errors can be judged by comparing the solution with the initial condition. Figure 36 below shows the fluid interface as it progresses through the simulation.




Figure 36 - Depiction of how the 3D sphere in a deformation field test case progresses through out the simulation. Starting most left, the sphere deforms from left to right. Halfway through the simulation, the flow field is reversed and the interface reforms into the initial condition going right to left.


Sphere in a Deformation Field Results: Non-Rotational Advection

A view of the fluid interface from the bottom of the deformation sphere at the half mark of the simulation is displayed in figure 37. The test case was tested on the following meshes: $N^3 = 64^3$, 128^3 , and 256^3 . A mesh sizing above 256^3 becomes very costly computationally.

When looking at figure 37, it can be seen that the expected trend is apparent when comparing finite difference, ELVIRA, and LVIRA. Firstly, the LVIRA method preserves the interface more completely than the finite difference and ELVIRA normal schemes. In the 64^3 mesh, the LVIRA method has the lowest amount of empty space at the bottom layer of the thinned sphere. Observing the mesh for 128^3 , it is also apparent that the LVIRA method provides the smoothest interface. The finite difference normal provides a wavy interface, and the ELVIRA normal also creates small ridges. This is compared to the LVIRA that provides the smoothest interface. When the mesh is refined to 256^3 , all the normal approximations are visually similar and are all smooth as expected.

When the shape error is plotted over the course of the simulation time a comparison between the different normal schemes can be quantified. Figure 38 on page 57, shows the plot of the shape error over the course of the simulation. From the figure 38 it is seen that generally the shape error for the LVIRA normal is the lowest. On the smallest mesh tested, figure 38a, after the half mark of the simulation, the ELVIRA normal error stops growing and begins to decrease along the same order as the LVIRA method. At the half mark, $t=1.5$, the LVIRA method performs the best, with the ELVIRA normal second and the finite difference normal showing the highest error of the three configurations tested.

Operator-Split Advection – Non-Rotational – $T = 1.5$
 Finite Difference ELVIRA LVIRA



N = 64



N = 128



N = 256

Figure 37 - Fluid interface halfway through the test case, when $t = 1.5$, after the sphere is initially deformed using the operator-split non-rotational advection scheme for transportation.

As the mesh becomes finer, $N=128$ and $N=256$, the ELVIRA method always shows the lowest shape error throughout the simulation. A general trend is that after the sphere starts to reform, the errors generally decrease. This is believed to be the cancellation of advection errors as was seen in the circle in a deformation field test case.

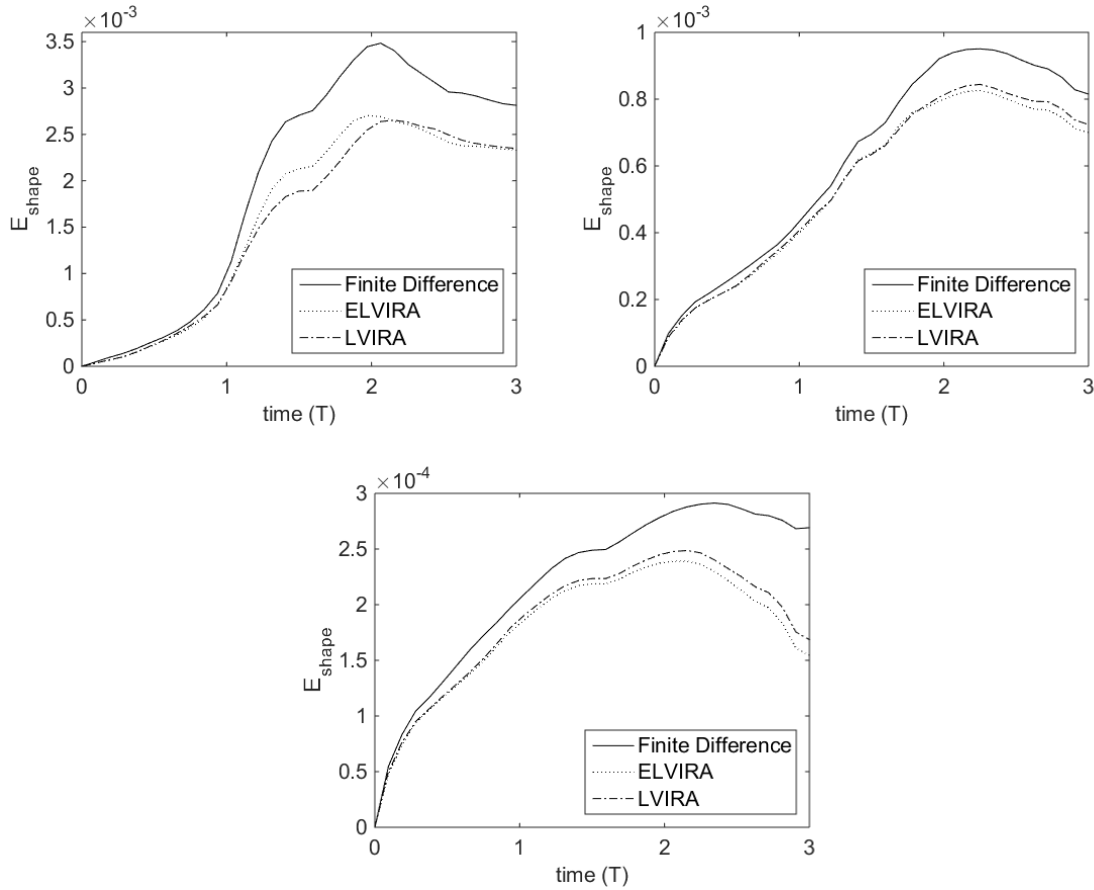


Figure 38 - Approximate shape errors plotted over the entire simulation. (a) Top left: $N=64$, (b) top right: $N=128$, (c) bottom: $N=256$.

The fluid interfaces at the end of the simulation are depicted on the next page in figure 39. From viewing the fluid interfaces some things are apparent. Firstly, in the coarser meshes there are apparent bubbles that have built up within the fluid interface. These bubbles are not a result of float-sum that has built up within the sphere. Instead they are a result of the breakdown in the thin interface when the sphere is completely stretched out. There is some visible float-sum that is visible within the sphere. These float-sum elements are harder to see, and are the small-disoriented planes that are spread throughout the domain, inside and outside of the sphere. One of the other observations that can be drawn is amount of float-sum that is created when using the ELVIRA normal

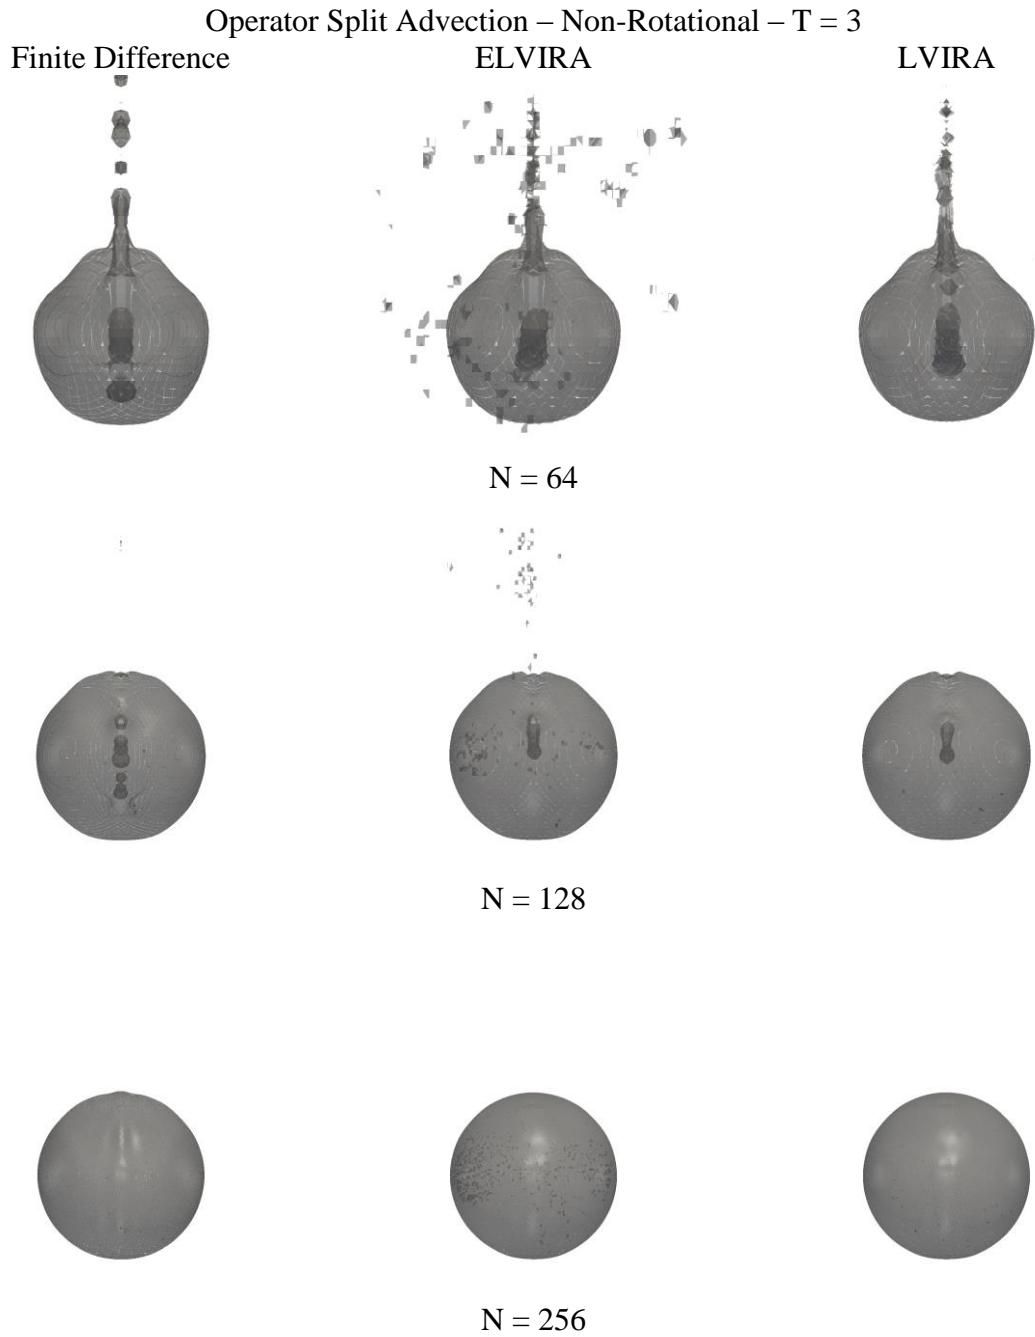


Figure 39 - Fluid interface at the end of the test case, when $t = 3$, after the sphere is initially deformed and reformed using the operator-split non-rotational advection scheme for transportation.

scheme. ELVIRA has a considerable amount when compared to the LVIRA and finite difference normal schemes.

The final shape errors for the operator split non-rotational advection scheme are displayed below in table 14. Figure 40 shows the convergence plot.




Figure 40 - Convergence history of the shape error in the deformation sphere test case using the operator-split non-rotational advection scheme.

Advection Scheme: Operator-Split Non-Rotational						
N	Finite Difference	Convergence Rate	ELVIRA	Convergence Rate	LVIRA	Convergence Rate
64	2.964E-03	-	2.559E-03	-	2.784E-03	-
128	8.150E-04	1.86	6.996E-04	1.87	7.237E-04	1.94
256	2.692E-04	1.60	1.543E-04	2.18	1.686E-04	2.10

Table 14 - Shape errors for the operator-split non-rotation advection scheme applied to the sphere in a deformation field test case.

The final shape errors show that the ELVIRA method shows the best convergence. When using the ELVIRA and LVIRA methods the operator split non-rotational advection scheme performs 2nd order in terms of shape error. The finite difference normal scheme performs near second order on the coarser meshes, but as the mesh is refined, the

convergence rate starts to drop below second order. At a mesh resolution of $N=256$, the ELVIRA and LVIRA methods perform very similar.

The timing of the PLIC operation is compared next. Figure 41 shows the comparison of the three normal geometric reconstruction schemes.




Figure 41 - Timing comparison of the Finite Difference, ELVIRA, and LVIRA normal schemes using the operator-split non-rotational advection scheme for the sphere in a deformation field test case.

Advection Scheme: Operator-Split Non-Rotational						
N	Finite Difference		ELVIRA		LVIRA	
	N_{PLIC}	Timing [ms]	N_{PLIC}	Timing [ms]	N_{PLIC}	Timing [ms]
64	30472	3.00	35418	65.34	50889	252.66
128	115756	17.19	123999	211.33	243573	951.22
256	513231	82.75	524347	837.85	1177519	4253.26

Table 15 - Average time per iteration spent during reconstruction of the PLIC interface using the operator-split non-rotational advection scheme. N_{PLIC} represents the amount of cells that are reconstructed in the band structure.

In terms of computational time spent per time-step to reconstruct the fluid interface, the finite difference was by far the fastest. Even on the coarsest mesh tested the PLIC reconstruction time was less than one-tenth of a second. While the ELVIRA took nearly a second to reconstruct per time-step on the finest mesh tested. The LVIRA method took

the longest to reconstruct on every mesh size tested. It should be noted that on the operator-split non-rotational advection scheme, there was an oddity that increased the number of cells that were reconstructed, N_{PLIC} . Nearly twice as many cells were triggered for reconstruction. This added to the overall time for LVIRA to reconstruct the interface.

Sphere in a Deformation Field Results: Rotational Advection

Since the flow field has a gradient along each principle axis direction, the split rotational scheme is also tested. Figure 42 shows the fluid interface from the bottom view at the half mark of the simulation when $T=1.5$. The interfaces are largely unchanged between the two operator splitting schemes. For the $N=64$ mesh, the finite difference method has the largest void for the stretched interface. The ELVIRA does better at maintaining the interface, and the LVIRA scheme shows the smallest void. There is still breakup however due to how thin the interface becomes with respect to the mesh size.

For the finer mesh of $N=128$, the finite difference shows wave like ripples in the surface, much like in the non-rotational operator split advection scheme. The ELVIRA method also shows the ridges that are also apparent in the non-rotational operator split advection scheme. And finally the LVIRA method is mostly smooth since it handles the thin interface the best. At the finest mesh, the visual differences are negligible.




The shape errors plotted as a function of simulation time are displayed on page 63, in figure 43. The plots are nearly identical to the operator split non-rotational scheme as well. In the $N=64$ mesh, the LVIRA shape error is the lowest until the sphere starts to reform, at which point the ELVIRA method dips below the LVIRA shape error. After the coarsest mesh, the advection scheme performs best when using the ELVIRA normal scheme. The LVIRA and ELVIRA normal schemes are very close to each other.

Operator Split Advection – Rotational – $T = 1.5$

Finite Difference

ELVIRA

LVIRA



$N = 64$



$N = 128$



$N = 256$

Figure 42 - Fluid interface halfway through the test case, when $t = 1.5$, after the sphere is initially deformed using the operator-split rotational advection scheme for transportation.

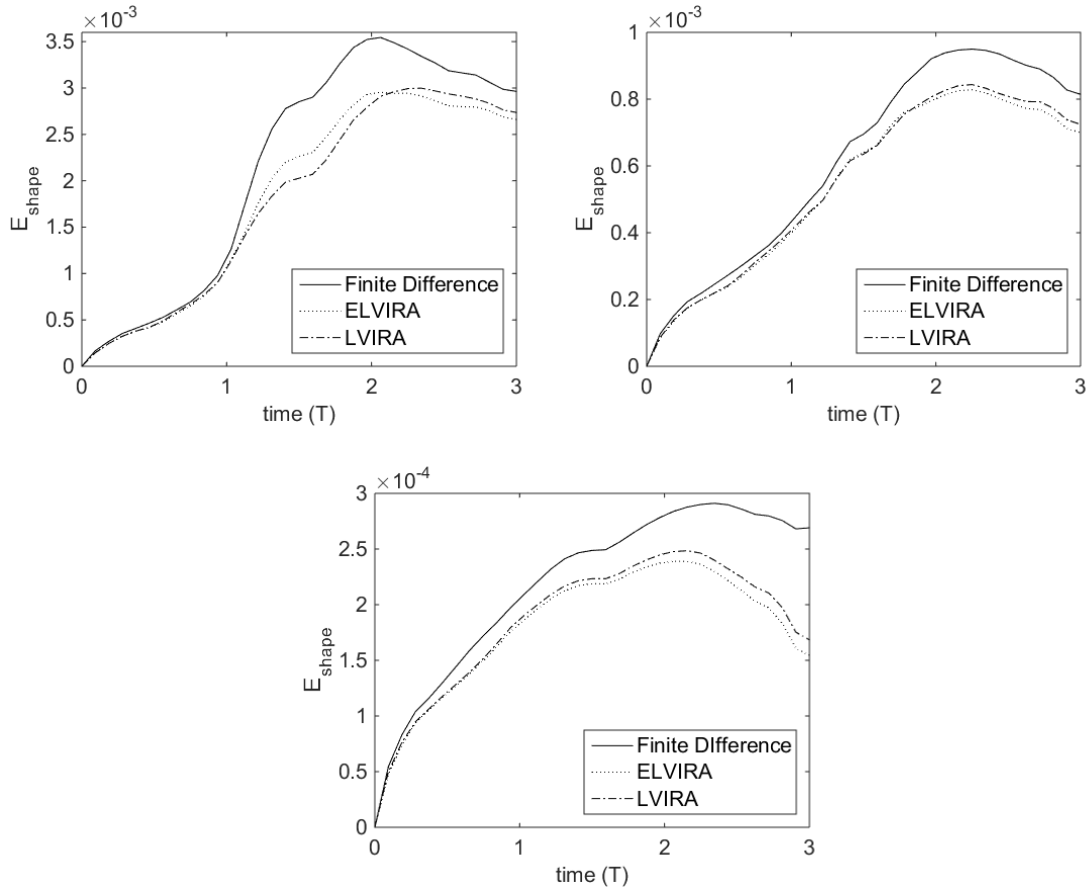


Figure 43 - Approximate shape errors plotted over the entire simulation. (a) Top left: $N=64$, (b) top right: $N=128$, (c) bottom: $N=256$.

The depictions of the fluid interface at the end of the simulation are showed in figure 44 on the next page. Again, many of the results are nearly identical to the operator split non-rotational advection scheme. The bubble within the sphere is still apparent. The amount of float-sum that is in the final results is much more than that when using the operator split non-rotational advection scheme. This is especially true for the ELVIRA normal scheme. The second column of figure 44 shows the results when using the ELVIRA normal scheme. The amount of float-sum is very high in the region right above the sphere.

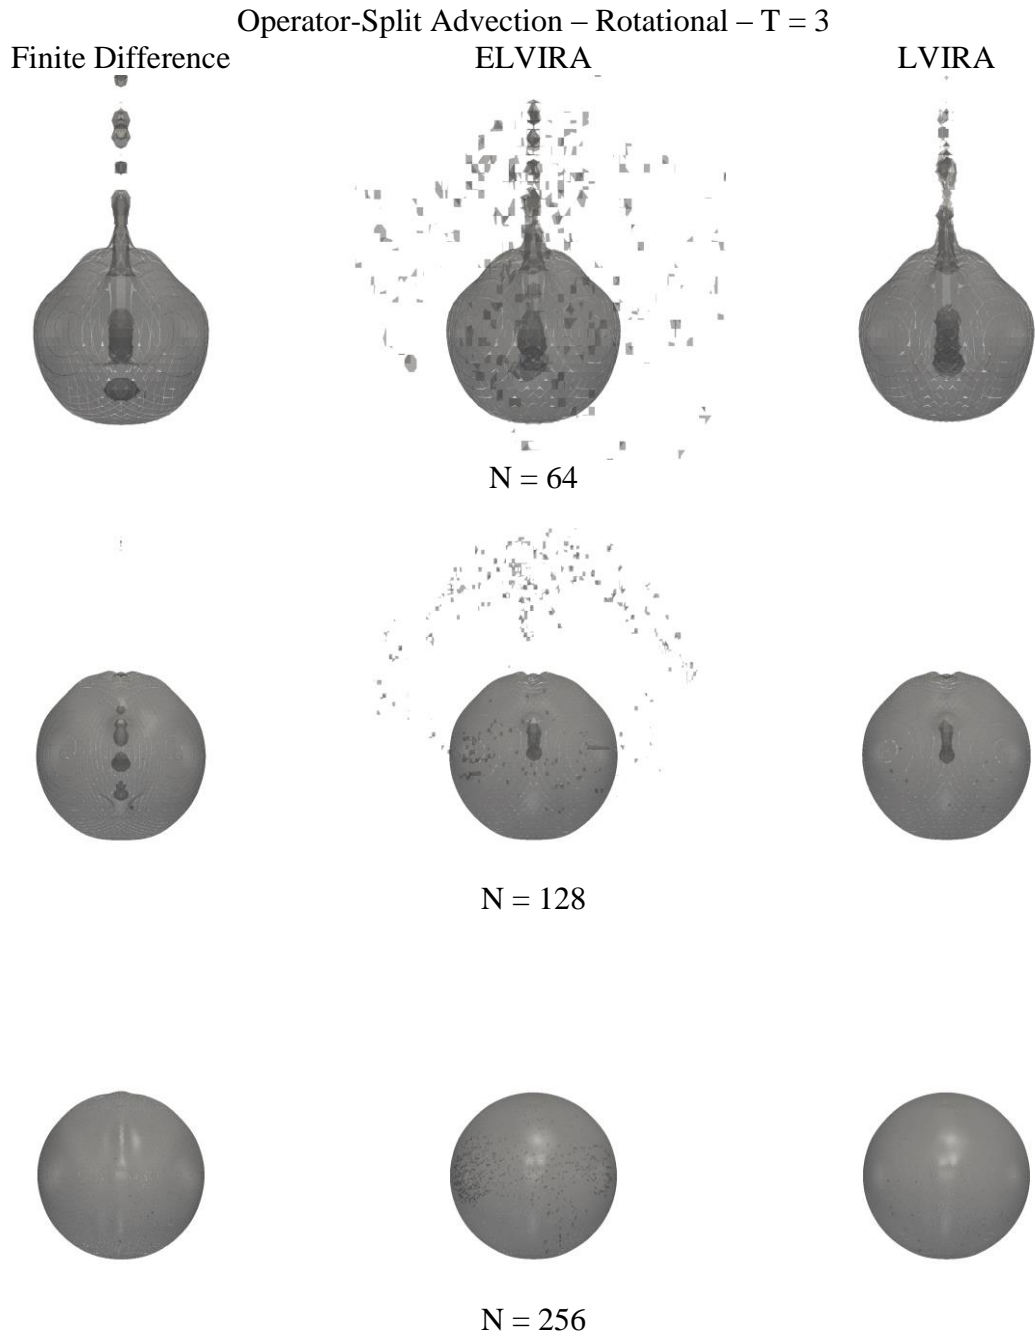


Figure 44 - Fluid interface at the end of the test case, when $t = 3$, after the sphere is initially deformed and reformed using the operator-split rotational advection scheme for transportation.

The final result shape errors are displayed next for the sphere in a deformation field test case when using the operator split rotational advection scheme.




Figure 45 - Convergence history of the shape error in the deformation sphere test case using the operator-split rotational advection scheme.

Advection Scheme: Operator-Split Rotational						
N	Finite Difference	Convergence Rate	ELVIRA	Convergence Rate	LVIRA	Convergence Rate
64	2.963E-03		2.660E-03		2.735E-03	
128	8.144E-04	1.86	6.998E-04	1.93	7.244E-04	1.92
256	2.691E-04	1.60	1.543E-04	2.18	1.684E-04	2.10

Table 16 - Shape errors for the operator-split rotational advection scheme applied to the sphere in a deformation field test case.

Inspection of the final shaper error also confirms that the added rotational component when transporting the PLIC interface does very little to affect the solution for the three-dimensional sphere in a deformation field test case. The convergence rate for the ELVIRA and LVIRA methods remains second order, and the finite difference method starts close to second order but loses this property on the finest mesh.

Timing data for the operator split rotational advection scheme is shown below. Figure 46 shows the comparison between the different normal schemes.




Figure 46 - Timing comparison of the Finite Difference, ELVIRA, and LVIRA normal schemes using the operator-split rotational advection scheme for the sphere in a deformation field test case.

Advection Scheme: Operator-Split Rotational						
N	N_{PLIC}	Finite Difference		ELVIRA		LVIRA
		N_{PLIC}	Timing [ms]	N_{PLIC}	Timing [ms]	N_{PLIC}
64	35022	42453	2.90	70.61	39781	238.30
128	162643	171630	20.44	250.61	214731	826.30
256	912064	925779	97.58	1123.53	960990	3359.02

Table 17 - Average time per iteration spent during reconstruction of the PLIC interface using the operator-split non-rotational advection scheme. N_{PLIC} represents the amount of cells that are reconstructed in the band structure.

Table 17, shows the removal of the oddity seen in the operator split non-rotational advection scheme where the LVIRA normal had nearly twice the cells that were being geometrically reconstructed. Looking at the N_{PLIC} columns, all the reconstructed elements are of the same order. This does not help the LVIRA method much as the LVIRA method still performs extremely slowly when compared to the other normal reconstruction schemes.

Sphere in a Deformation Field Results: Unsplit Advection

The geometric unsplit advection scheme is the last advection scheme that was applied to the test case of the sphere in a deformation field test case. The fluid interface is displayed at $t=1.5$ for each of the mesh refinement levels in figure 47. At each mesh sizing, both of the operator-split advection schemes were comparable to the unsplit advection scheme. Overall, the sizing of the void within the thin surface of the fluid interface remains constant in the coarsest mesh, $N=64$, despite using the different advection schemes. The largest differences can be seen by using better normal approximation schemes. By using the LVIRA normal scheme, the PLIC interface is better maintained for accurately calculating fluxes.

In the $N=64$ test case, the LVIRA normal as expected preserves the interface the most when compared to the EVIRA and finite difference normal schemes. The ELVIRA is next best at preserving the interface, but has more regions where the PLIC normal artificially breaks apart. The finite difference performs the worst.

In the $N=128$ mesh sizing, the interfaces show nearly identical to the operator-split advection schemes. The finite difference method has least smooth interface. Many of the wave features are in the same positions as the operator-split advection schemes. The EVLIRA fluid interface also contains the small ridges that were seen when using the operator split advection schemes. The ridges are less pronounced than the wave features seen in the finite difference normal interface. Lastly, the LVIRA method continues to show its ability to preserve a smooth interface.

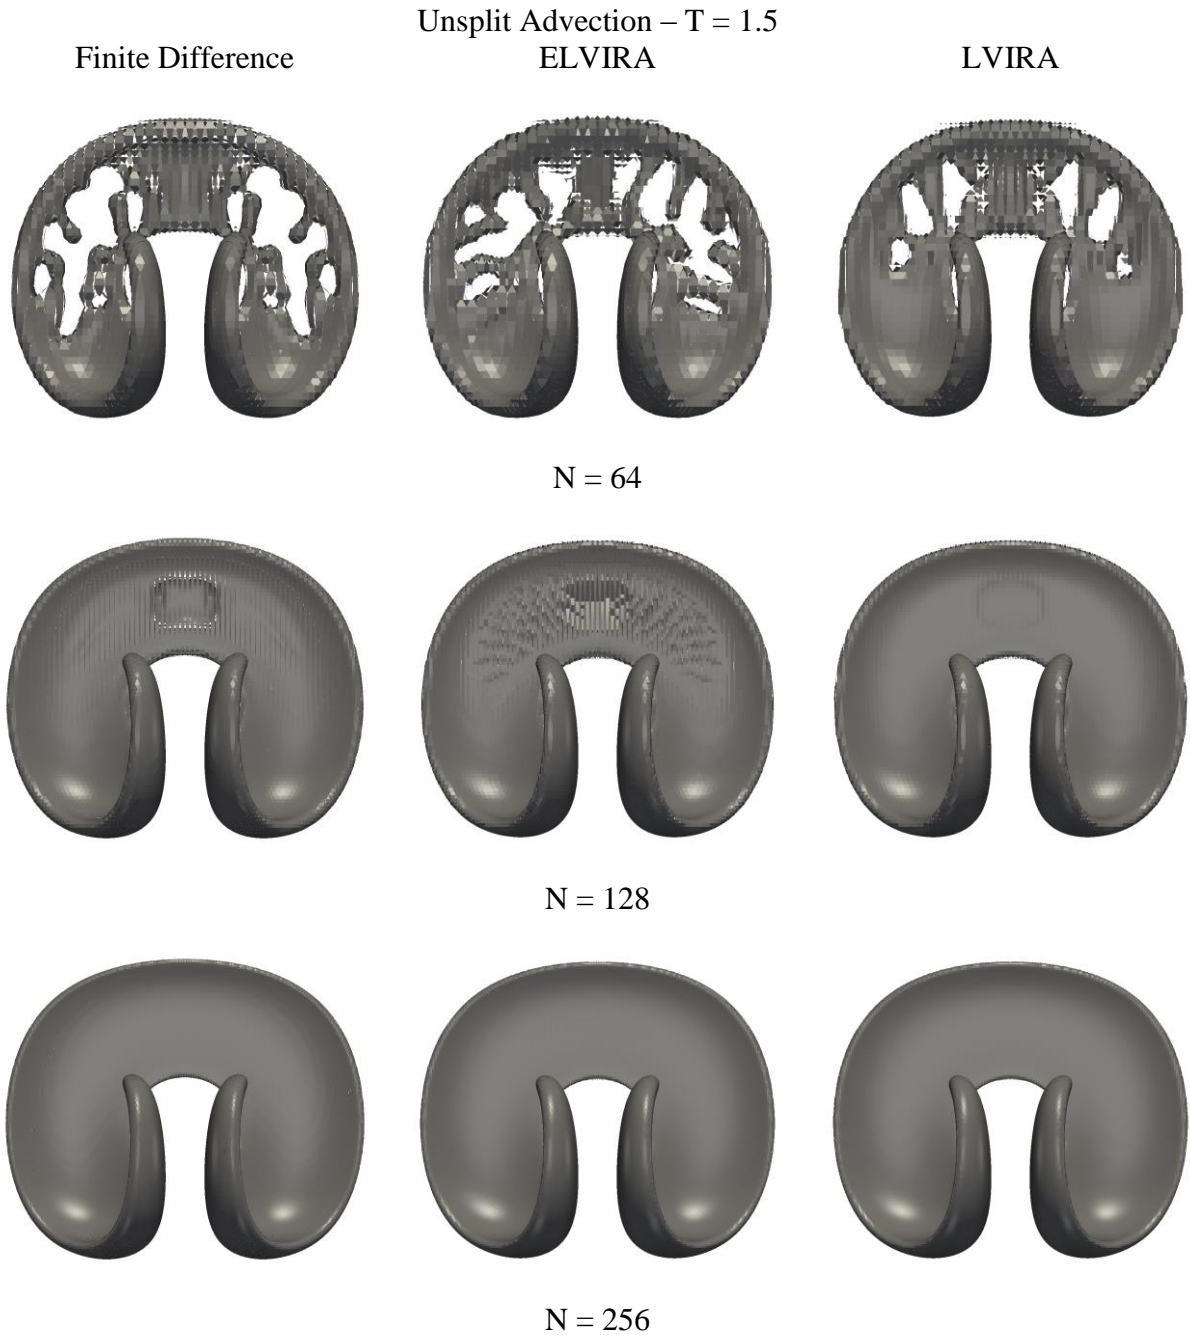


Figure 47 - Fluid interface halfway through the test case, when $t = 1.5$, after the sphere is initially deformed using the unsplit advection scheme for transportation.

Lastly, in the finest mesh, $N=256$, all surfaces are smoothed out and no differences are apparent. It is nearly impossible to distinguish any differences from the choice of advection scheme.

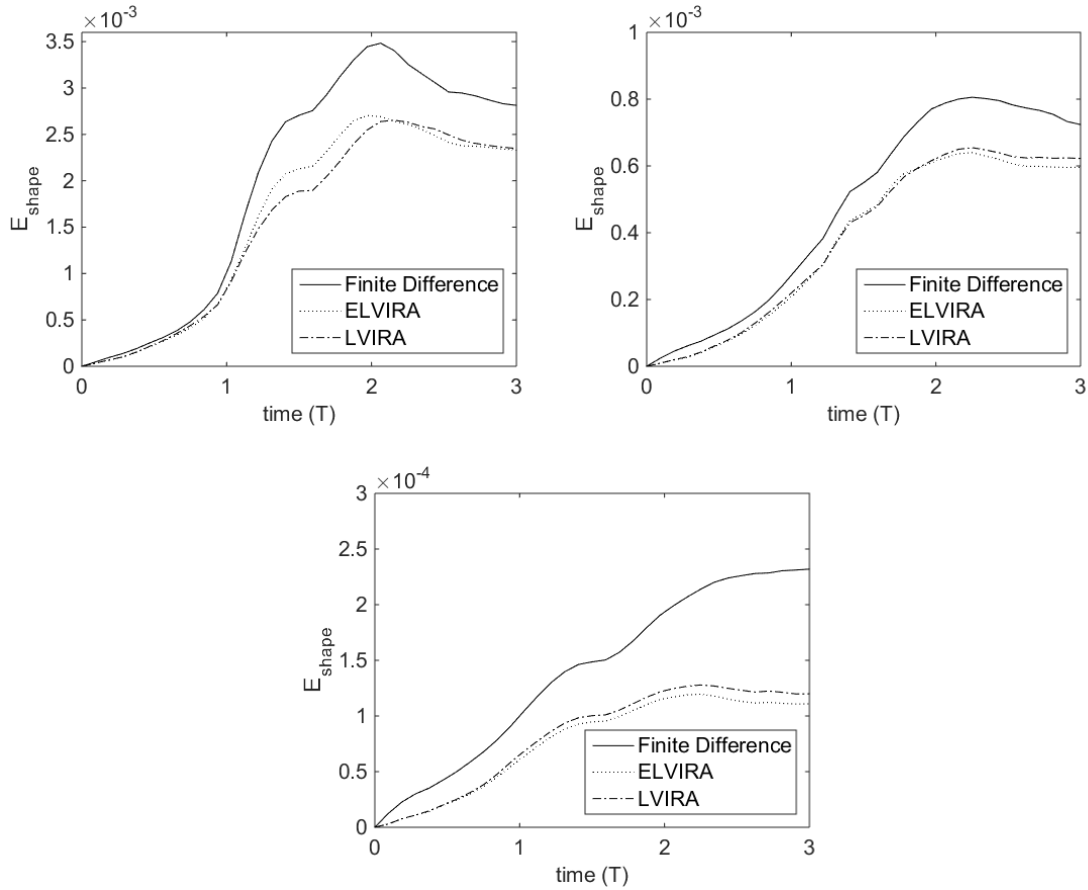


Figure 48 - Approximate shape errors plotted over the entire simulation. (a) Top left: $N=64$, (b) top right: $N=128$, (c) bottom: $N=256$.

The approximate shape errors plotted during simulation give better insight to how the unsplit method and each of the normal comparisons perform. First, the curves in figure 48 follow the same trend as seen in the operator-split advection methods. For $N=64$, LVIRA generally shows the lowest shape error until the sphere starts to reform at which point the shape error using the ELVIRA normal drops below the LVIRA shape error. For $N=128$ and 256 meshes, the curves are the same shape as the operator-split methods, but the scales are different. In the case of the finest mesh tested, the shape errors for the ELVIRA and LVIRA normal are close to $1e-04$ at a time of 1.5. This is more than half the value of the shape error when using either of the operator split advection schemes.




Figure 49 - Fluid interface at the end of the test case, when $t = 3$, after the sphere is initially deformed and reformed using the unsplit advection scheme for transportation.

The final shapes of the fluid interface are shown above in figure 49 when the unsplit advection scheme is used for transportation in the sphere in a deformation field test case.

The shapes are generally the same as before in the operator-split test cases. On the coarsest mesh tested the bubble within the reformed sphere is visible. This confirms the

bubble is an artifact of the shape error and not due to float-sum formation. However, on the finest mesh that was tested, 256^3 , there is float-sum that has formed. Multiple times during the course of this work, when the mesh was sufficiently refined float-sum would begin to form despite using the conservative unsplit geometric transportation scheme. This is believed to be due to round-off error when transporting fluid volumes. When one examines a fully immersed cell that is far enough from the interface (such that it remains fully immersed after one transportation step), the divergence free flow field requires that the net flux through every face of the cell be equal to zero. Due to round-off errors, this requirement cannot be exactly met. After transportation, instead of having a volume fraction of 1 (fully immersed), round-off errors cause the volume fraction to be 0.99999999 or 1.00000001. When round-off errors allow the volume fraction to fall below a reconstruction threshold, C_{full} , float-sum forms. This leads to an argument that for the volume fraction to remain a physical value, the volume fraction field must be cleaned up after transportation. This is done by setting any volume fraction that is above, C_{full} to a volume fraction of exactly one. The same is done for the opposite end of the volume fraction, but instead of a C_{full} value, a C_{empty} threshold is defined. Any volume fraction below C_{empty} is set to exactly zero. Unfortunately, this affects the conservation of volume properties of the volume-of-fluid method. If C_{full} is set close enough to one (say $1 - 1e-12$) then these added volume fractions do not greatly affect the solution and are of the same order as the round-off error.




Figure 50 - Convergence history of the shape error in the deformation sphere test case using the operator-split rotational advection scheme.

Advection Scheme: Unsplit						
N	Finite Difference	Convergence Rate	ELVIRA	Convergence Rate	LVIRA	Convergence Rate
64	2.813E-03		2.333E-03		2.348E-03	
128	7.238E-04	1.96	5.958E-04	1.97	6.224E-04	1.92
256	2.319E-04	1.64	1.109E-04	2.43	1.199E-04	2.38

Table 18 - Shape errors for the operator-split rotational advection scheme applied to the sphere in a deformation field test case.

The final shape errors in table 18 show slight decreases in shape error from using the unsplit advection scheme. The percentage drop from switching from the operator-split advection scheme to the unsplit scheme yields 13.79%, 28.11%, and 28.82% for the finite difference, ELVIRA, and LVIRA normal schemes respectively, on the finest mesh tested. Generally, the highest drops were seen from using the ELVIRA or LIVRA normal approximation scheme. The convergence rates were also slightly higher for the unsplit

schemes as the convergence rates for the N=256 mesh were close to 2.4 for the ELVIRA and LVIRA normal approximations.




Figure 51 - Timing comparison of the Finite Difference, ELVIRA, and LVIRA normal schemes using the unsplit advection scheme for the sphere in a deformation field test case.

Advection Scheme: Unsplit						
N	Finite Difference		ELVIRA		LVIRA	
	N_{PLIC}	Timing [ms]	N_{PLIC}	Timing [ms]	N_{PLIC}	Timing [ms]
64	24177	1.04	25459	16.31	27138	59.95
128	93687	5.05	93858	48.41	111309	145.17
256	417861	25.33	429167	188.57	440527	512.86

Table 19 - Average time per iteration spent during reconstruction of the PLIC interface using the operator-split non-rotational advection scheme. N_{PLIC} represents the amount of cells that are reconstructed in the band structure.

The comparison of the average PLIC reconstruction timing per time-step is shown in figure 51. The time spent in PLIC reconstruction is significantly reduced when using the unsplit geometric advection scheme due to the removal of float-sum in the domain. The removal of float-sum reduces the number of cells that the reconstruction algorithm targets for reconstruction. In comparison to the finite difference and ELVIRA normal

approximation schemes, the LVIRA method is the highest costing. The finite difference scheme produces the fastest result due to its simplicity.

Sphere in a Deformation Field Results: Advection Timing Comparisons




Figure 52 – Comparison of average time spent during the advection step of the volume-of-fluid algorithm for the sphere in a deformation field test case.

N	Advection Timing [ms]		
	Split Non-Rotational	Split Rotational	Unsplit
64	2.78	2.57	400.31
128	21.90	22.49	1566.75
256	156.16	137.53	6437.04

Table 20 – Average computational time per iteration for each advection scheme tested in the sphere in a deformation field test case.

In three-dimensions, the unsplit method becomes extremely costly. The time spent in the transportation of the volume fraction is anywhere from 40 to 150 times as high when using the geometric unsplit advection method. This would be acceptable if the errors were 40 times smaller, but this was not the case for the sphere in a deformation field test case.

CHAPTER 6

CONCLUSION

In conclusion, three different advection algorithms and three difference transportation algorithms have been compared to see the effect on errors and computational timing.

In the first test case of observing the convergence of the three normal approximations, it was found that the Efficient Least-Squares Volume-of-Fluid Interface Reconstruction Algorithm and Least-Squares Volume-of-Fluid Interface Reconstruction Algorithm both produced a normal that was 1st order accurate. While this is acceptable, a 2nd order approximation of the interface normal would be preferable since the best geometric fluxing approximation is 2nd order accurate. In contrast, the finite difference normal approximation was found to be 0th order accurate since it would converge asymptotically and eventually be unable to approximate the normal direction sufficiently.

In the second test case of the Zalesak disk, the normal comparisons showed nearly no differences. ELVIRA and LVIRA had a slight advantage, but no large difference were apparent. This was also the case when using the three different types of transportation algorithms. All methods converged better than 1st order, except the finite difference which was only slightly better than 1st order.

In the third test case of the circle in the deformation field test case, there were many differences when using the three different normal approximation schemes and three different transportation algorithms. When using the three different normal approximation schemes, the LVIRA normal calculation showed the best at preserving a thin interface, while the ELVIRA normal calculation showed the second best at this trait, and the finite

difference approximation showed the worst ability to maintain a thin interface. Unfortunately these differences in preserving the interface were only slightly apparent in the final shape error calculation of the interface. This was believed to be due to the time reversal of the flow field, where there is error that cancels out when the ligament is reformed. The geometric unsplit method also showed improvements in the transportation of the interface. These were more apparent in the shape error calculations. The geometric unsplit method eliminated float-sum which is desirable for accurate interface capturing, since it maintains the volume fraction to physically meaningful values of $0 \leq C \leq 1$. In terms of timing, it was shown that the LVIRA method took the longest time to approximate the normal since it uses an iterative procedure to determine the appropriate normal. The Efficient-LVIRA method took less time to compute the normal, but in terms of shape error the ELVIRA and LVIRA methods performed the same (keep in mind that LVIRA still maintained the interface better). It was also shown that the unsplit geometric algorithm took the longest time to compute fluxes, while both operator-split fluxing operations took significantly less time. It is this test case that best demonstrates the differences between the normal and advection schemes in this work. Figure 53, on the next page, shows a plot of how the shape error varies as a function of computational cost for the methods tested in this work. A line connects each of the combinations of normal calculation and advection scheme for the various mesh sizes tested. From figure 53, it can be seen that the unsplit advection uses the most computational time, but shows the lowest shape error. It can also be seen that when using the unsplit advection scheme, the shape error has a steeper slope as compared to the operator-split advection schemes. The increase in slope shows the unsplit advection scheme is advantageous for shape errors

below $1e-03$. The most apparent reason for this is the removal of float-sum when using an unsplit advection scheme. The addition of float-sum that is a result of the operator-split advection scheme increases the cost due to an increased number of elements that must be reconstructed and transported. When the problem size becomes large enough, the number float-sum elements become so large that the unsplit method performs faster as compared to the operator-split methods.




Figure 53 – Log-Log plot of Shape Error vs Average Computation Time per Timestep for all combinations of advection and normal calculation schemes used in the Deformation Column test case.

Finally, in the last test case, the sphere in a deformation field test case, we saw differences similar to those seen in the circle in a deformation field test case. The LVIRA normal was by far the best at maintaining a smooth, thin interface. The ELVIRA method was second in this ability, and the finite difference method was last. Final shape errors showed little differences in the LVIRA and ELVIRA normal approximations, but more apparent differences when compared to the finite difference normal approximation. Transportations methods showed little differences in this test case and both ELVIRA and LVIRA maintained higher than 2nd order accuracy regardless of chosen transportation algorithm. Timing also showed similarity to the circle in a deformation field test case in that the finite difference method was the lowest costing, the ELVIRA was second, and LVIRA was most costly. In terms of advection both operator-split methods took approximately the same time, while the unsplit advection algorithm was the most costly.

It is still difficult to tell clearly which method is the best, the LVIRA normal is best at maintaining thin interfaces on course meshes. However on very refined mesh sizes, the ELVIRA and LVIRA method perform the same, while the LVIRA normal costs much more than the ELVIRA method computationally. One would have to anticipate how important the interactions of small scale structures, like thin interfaces, are when modeling a given problem. If thin ligaments between fluid interfaces become important, the LVIRA method will be preferable, but if interfaces remain thick, the ELVIRA is more preferable due to the decrease in computational cost.

In terms of transportation algorithm, the 2nd order unsplit advection scheme shows the most stability in maintaining 2nd order accuracy. While the operator-split advection

methods should be 2nd order accurate (Pilliod & Puckett, 2004), they showed fluctuation, jumping between 1st and 2nd order.

REFERENCES

- Gueyffier, Denis, Jie Li, Ali Nadim, Ruben Scardovelli, and Stéphane Zaleski. "Volume-of-Fluid Interface Tracking with Smoothed Surface Stress Methods for Three-Dimensional Flows." *Journal of Computational Physics* 152.2 (1999): 423-56. Print.
- Li, Jie. "Calcul d'interface affine par morceaux (piecewise linear interface calculation). *CR Acad. Sci. Paris, Sér. Lib* 320 (1995): 391-96. Print.
- Noh, W. F., and Paul Woodward. "SLIC (Simple Line Interface Calculation)." *Proceedings of the Fifth International Conference on Numerical Methods in Fluid Dynamics* (1976): 330-40. Print.
- Owkes, Mark, and Olivier Desjardins. "A Computational Framework for Conservative, Three-dimensional, Unsplit, Geometric Transport with Application to the Volume-of-Fluid (VOF) Method." *Journal of Computational Physics* 270 (2014): 587-612. Print.
- Parker, B.J., and D.L. Youngs. Two and Three Dimensional Eulerian Simulation of Fluid Flow with Material Interfaces. Aldermaston, Berkshire: Atomic Weapons Establishment, 1992. Print.
- Pilliod, James Edward, and Elbridge Gerry Puckett. "Second-order Accurate Volume-of-fluid Algorithms for Tracking Material Interfaces." *Journal of Computational Physics* 199.2 (2004): 465-502. Print.
- Puckett, Elbridge G. "A Volume-of-fluid Interface Tracking Algorithm with Applications to Computing Shock Wave Refraction." *Proceedings of the Fourth International Symposium on Computational Fluid Dynamics* (1991): 933-38. Print.
- Renardy, Yuriko, and Michael Renardy. "PROST: A Parabolic Reconstruction of Surface Tension for the Volume-of-Fluid Method." *Journal of Computational Physics* 183.2 (2002): 400-21. Print.
- Rider, William J., and Douglas B. Kothe. "Reconstructing Volume Tracking." *Journal of Computational Physics* 141.2 (1998): 112-52. Print.
- Scardovelli, Ruben, and Stéphane Zaleski. "Direct Numerical Simulation of Free-Surface and Interfacial Flow." *Annual Review of Fluid Mechanics* 31.1 (1999): 567-603. Print.