

Feature-Driven Cartesian Adaptive Mesh Refinement in the Helios Code

Sean J. Kamkar *

Andrew M. Wissink ‡

Antony Jameson †

Venkateswaran Sankaran §

Stanford University, Palo Alto, CA 94305

Ames Research Center, Moffett Field, CA 94035

Feature-detection methods are applied to drive Cartesian-based adaptive mesh refinement for improved vortex resolution in aerodynamics flowfields. Several approaches such as the Q -criterion, the eigenvalues of velocity gradient tensor and the identification of local pressure minima are considered. Specific attention is given to the automation of the refinement process so that it is applicable to a wide range of flowfields without need for user intervention. This is achieved by careful normalization of the methods and by applying appropriate threshold values to identify specific regions for mesh refinement. The methods are implemented within the Helios code, which features a dual-mesh paradigm based on using unstructured meshes in the near-body region and adaptive Cartesian meshes in the off-body region. The resultant adaptive code is evaluated for both theoretical and practical test problems, and the results are compared with those obtained by uniform mesh resolution. It is demonstrated that the adaptive solutions provide comparable accuracy for a fraction of the cost of using uniform meshes and thus show promise for efficient and accurate simulations of vortex-dominated aerodynamics flows.

I. Introduction

It is well known that the accuracy of computational fluid dynamics (CFD) simulations can be significantly improved by increasing the number of grid points. However, it is impractical to uniformly improve the mesh resolution throughout the entire computational domain since a large proportion of grid points would be wasted in regions that are not relevant to enhancing the accuracy. Adaptive mesh refinement (AMR)

*Graduate Student, Durand Building, Department of Aeronautics & Astronautics, skamkar@stanford.edu, AIAA Member

†Professor, Durand Building, Department of Aeronautics & Astronautics, jameson@baboon.stanford.edu, AIAA Fellow

‡Aerospace Engineer, US Army Aeroflightdynamics Directorate, andrew.m.wissink@us.army.mil, AIAA Member

§Aerospace Engineer, US Army Aeroflightdynamics Directorate, vsankaran@merlin.arc.nasa.gov, AIAA Member

schemes provide a means of placing points only in those regions where they are necessary and can thereby enhance accuracy in an efficient manner. With the dramatic increases in available computational resources, along with the concomitant demands for improved accuracy, AMR schemes are poised to become increasingly important. The focus of the present paper is on developing schemes that are suitable for Cartesian-based AMR and targeted specifically at vortex-dominated flows.

Rotorcraft flow-fields present unique challenges to computational aerodynamics. Primary among these is the need to accurately capture and preserve the trailing tip vortices that are shed from the spinning rotor blades. The vortices can subsequently interact with the blades themselves, the fuselage, and the tail rotors. A further important aspect of rotorcraft problems is the close coupling between the CFD and the computational structural dynamics (CSD) solutions. In order to accommodate these challenges, the US Army's High Performance Computing Institute for Advanced Rotorcraft Modeling and Simulations (HI-ARMS) is developing a modular CFD/CSD comprehensive analysis package called Helios.¹⁻³ The CFD component of Helios uses a dual-mesh solution method that is composed of an unstructured solution strategy in the near-body domain, coupled using overset techniques with a high-order adaptive Cartesian framework in the off-body domain. The feature detection study, that is the subject of the present paper, is primarily motivated by the need to automate the Helios framework and to provide adaptive Cartesian refinement that specifically targets vortical features in the off-body region.

Pioneering work on hierarchical structured Cartesian-grid AMR was performed by Berger and Olinger⁴ and Berger and Colella.⁵ This approach was extended to structured curvilinear grids by Berger and Jameson.⁶ Error in the solution was used to determine regions for refinement. Their research highlighted the effectiveness of AMR and demonstrated that considerable gains in efficiency and accuracy could be obtained. Recent AMR work has focused on two possible paths for reducing solution error: adjoint-based analysis⁷⁻⁹ and feature-based analysis. Although adjoint-based methods have some attractive aspects, they are presently limited in their ability to handle unsteady problems such as those encountered in rotorcraft dynamics. Moreover, as discussed earlier, rotorcraft flow-fields are characterized by the need to resolve the blade tip vortices, which are ideally suited for feature-based strategies.

Previous work by the present authors focused on the design and development of feature-detection techniques to guide the mesh refinement process for vortex-dominated flows.¹⁰ A critical component in the design of these methods was the automation of the refinement process so that a wide range of flow scales could be accommodated without the need for user intervention. Traditional approaches such as the vorticity magnitude or the Q -criterion are unsuitable in this regard because of the problem-dependent nature of the threshold values of these parameters that are needed to control the regions that are tagged for refinement. Specifically, four feature detection methods were proposed, including the Q -criterion, pressure minimum sensors, eigen-

values of the velocity gradient tensor, and the correlation between the symmetric and anti-symmetric parts of the velocity gradient tensor. In all cases, appropriate normalizations and threshold values were identified to enable automated tagging of cells for refinement. These methods were evaluated using static solutions for a series of canonical vortex flows, and key parameters such as tagging accuracy and efficiency were catalogued. Importantly, these studies indicated that all of the non-dimensional feature-detection methods were comparable in terms of efficiency and accuracy and showed excellent promise for mesh adaption of vortex-dominated flows.

The objective of the present paper is to implement the feature-detection techniques within the Cartesian off-body framework of the Helios code and to systematically evaluate their performance for dynamically adapted problems. The implementation is carried out within the Guided Adaptive Mesh Refinement (GAMR) toolbox, which is a separate module that is used as a library by the off-body Cartesian solver in the Helios code. The combined adaptive-mesh code is subsequently evaluated for both theoretical and practical problems, such as the Lamb vortex, multiple interacting vortices, a three-dimensional ring vortex, and the NACA0015 wing. In all cases, performance estimates are made in terms of accuracy and efficiency, and the results are compared with those obtained with uniform fine mesh resolution.

The remainder of the paper is organized as follows. Section II describes the dual-mesh solution paradigm used in the Helios code. Section III presents the formulation of general feature detection algorithms that are implemented within the Cartesian infrastructure of Helios. In Section IV, we examine the performance of the feature detection schemes for a suite of theoretical and practical vortex-dominated flows. After this, we conclude by highlighting the strengths of the proposed approach and outlining some directions for future research and development.

II. Computational Infrastructure

Before presenting our results, a brief overview of the CFD infrastructure will help clarify the current implementation. We employ the Helios platform, which utilizes a dual-mesh overset approach—i.e, unstructured mixed-element grids near the body surface and structured Cartesian grids away from the surface.³ The rationale of this mixed near/off-body meshing scheme is to use unstructured grids near the surface to resolve complex geometry and boundary layer effects and to apply structured Cartesian grids away from the wall to accurately resolve the vortices in the far-field wake using time-dependent adaptive refinement (Figure 1).

The near-body unstructured solver is the NSU3D code¹¹ that uses second-order spatial discretization and implicit second-order backward differences in time for the solution of the unsteady Navier-Stokes equations. The off-body calculations are handled by the SAMARC component, which is an adaptive Cartesian code that combines an Euler CFD solution by the ARC3DC code developed by Pulliam¹² and a block-structured AMR

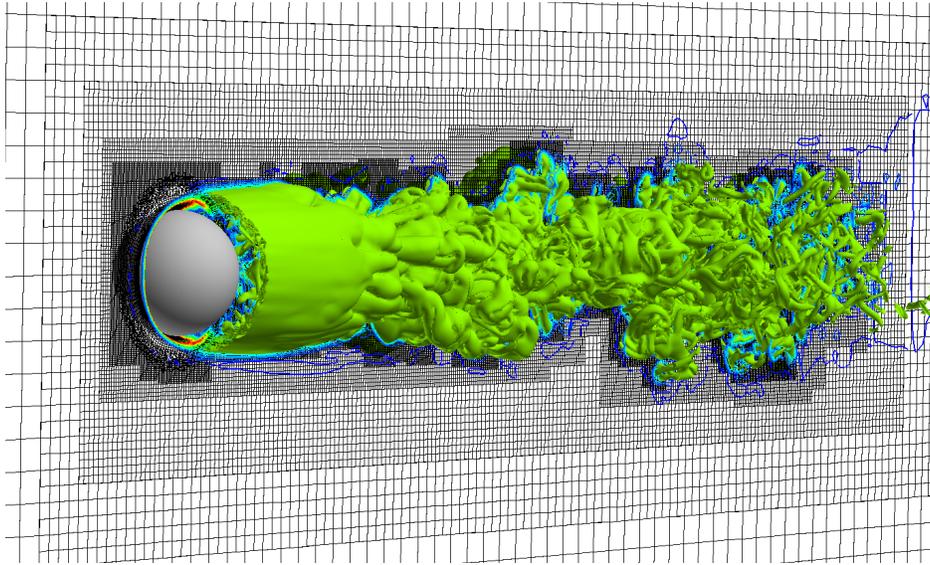


Figure 1: Isosurface of vorticity for flow over a sphere $Re = 800$ using near/off-body gridding approach.

scheme managed by SAMRAI.^{13–15} ARC3DC uses high-order algorithms optimized for Cartesian grids and calculates the flow solution on each Cartesian block. In this work, we apply 5th and 3rd order accurate spatial and temporal discretizations respectively. SAMRAI manages the parallel AMR operations, the construction and adaptation of the grid system, as well as parallel load balancing and MPI-based data exchanges between blocks. The Cartesian grid system is stored as a Berger and Colella⁵-style multi-level block-structured AMR (SAMR) grid hierarchy. Grid levels are constructed from coarsest to finest. The coarsest level defines the physical extent of the computational domain and each finer level is formed by selecting cells on the coarser level and then clustering the marked cells together to form block regions that will constitute the new finer level. The result is a hierarchy of nested refinement levels with each level composed of a union of logically-rectangular grid regions.

The domain connectivity between the near- and off-body grids is handled by PUNDIT, developed by Sitaraman *et al.*¹⁶ PUNDIT performs two main tasks: (1) Chimera grid hole-cutting to “iblack” points in the mesh that will not compute a solution as well as to identify points that exchange data, i.e. the inter-grid boundary points (IGBPs), and (2) the actual interpolation and inter-processor exchange of solution data between the solvers. It uses a 2nd-order interpolation algorithm and exchanges parallel data using MPI.

Integration of the three packages (NSU3D, SAMARC, and PUNDIT) is performed through a Python-based infrastructure.¹ The specific advantage of this approach is that it allows each code module to be treated as an object, providing a convenient way to assemble a complex multi-disciplinary simulation in an object-oriented fashion. Data exchanges are done without memory copies or file I/O, and the infrastructure can be run on large multi-processor clusters.¹⁷ As long as the Python interfaces are used only at the high-level, i.e.,

for managing data and not for computing numerics, the associated overheads are minimal.

III. Feature Detection Methods

Four feature detection methods are presented in this section, which are based on (1) the Q -criterion,¹⁸ (2) the λ_2 criterion,¹⁹ (3) the eigenvalues of the velocity gradient tensor $\nabla\mathbf{u}$,²⁰ and (4) the correlation between the symmetric and antisymmetric parts of $\nabla\mathbf{u}$.²¹ All methods are presented in a non-dimensional form by imposing a local normalization, which emphasizes generality and automation. For each method, we present some theoretical background and define a threshold function, $f_{threshold}$, which is used for purposes of vortex detection. This function is evaluated in each cell and the cell is marked for refinement if the resulting functional value is greater than a pre-specified value of the threshold, i.e., the cell is marked for refinement if $f_{threshold} > t_{val}$.

III.A. Non-dimensional Q

We begin by considering the rotation rate tensor (Ω), which is defined as the anti-symmetric part of $\nabla\mathbf{u}$ and represents local flow rotation. By applying the Frobenius norm, which for a real matrix is $\|M\| = [\text{trace}(MM^T)]^{\frac{1}{2}}$, we can estimate tensoral strength. This definition elucidates the significance of $\|\Omega\|$, since $2\|\Omega\|^2 = \|\boldsymbol{\omega}\|^2$, where $\boldsymbol{\omega}$ represents vorticity. In addition to considering Ω , we can also compute the strain rate tensor (S), which is the symmetric part of $\nabla\mathbf{u}$ and represents the local strain rate. Using both parameters we can obtain a measure of the *relative* vortical strength. Hunt¹⁸ was one of the first to assess the importance of this relationship, which is often referred to as Q and is usually defined as:

$$Q = \frac{1}{2} (\|\Omega\|^2 - \|S\|^2). \quad (1)$$

Beyond quantifying the difference between vorticity and strain rate, Q has additional physical interpretations. It acts as a source term in the Poisson equation for pressure $\nabla^2 p = 2\rho Q$, which is derived from the incompressible form of the Navier-Stokes equations (neglecting time dependent phenomena, and body and viscous forces).¹⁹ Q also appears as the second invariant in the characteristic equation for the velocity gradient tensor ($\nabla\mathbf{u}$), which is further explained in Section III.C.

Although Equation 1 measures the difference between the rotation and strain rate magnitudes, this value is still dependent upon the characteristic length and velocity scales of the problem. To yield a suitable non-dimensional form, we divide by $\|S\|^2$ to obtain a threshold function of:

$$f_{threshold} = \frac{1}{2} \left(\frac{\|\Omega\|^2}{\|S\|^2} - 1 \right). \quad (2)$$

Irrotational flow occurs when $f_{threshold} \rightarrow -\frac{1}{2}$, and solid body rotation happens when $f_{threshold} \rightarrow \infty$. Although $f_{threshold}$ is unbounded when $\|S\| \rightarrow 0$, this is acceptable since tagging will routinely occur when the function exceeds a constant, finite threshold value. Positive thresholds indicate regions where the vorticity strength is larger than the shear rate strength. A similar non-dimensional form was explored by Truesdell²² who formulated the non-dimensional parameter $N_k = \|\Omega\|^2/\|S\|^2$, which regardless of flow speed, determines the “quality” of rotation.

III.B. Non-dimensional λ_2

A region of low pressure typically occurs within a region of vortical motion. In the case of steady inviscid planar flow, a pressure minimum will occur along the axis of swirling motion because the pressure force must balance the centrifugal force. Such a force equilibrium is commonly referred to as a cyclostrophic balance. However, it has been shown¹⁹ that an unsteady strain rate may cause a pressure minimum in unsteady irrotational flow, which indicates that using only a pressure minimum to describe a region of rotational flow is not adequate. Despite this fact, Jeong¹⁹ *et. al.* have used this as a starting point for vortex identification. Discarding body forces, the vector form of the incompressible Navier-Stokes Equations is

$$\left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u}\right) = \frac{1}{\rho} (-\nabla p + \mu \nabla^2 \mathbf{u}), \quad (3)$$

where \mathbf{u} is velocity and μ is viscosity. We can take the gradient of the entire equation, which will allow us to refer to the term on the left hand side as the acceleration gradient $a_{i,j}$. This acceleration gradient can be decomposed into the same symmetric (S) and antisymmetric (Ω) parts as before. After neglecting unsteady strain and viscosity, and rewriting¹⁹ Equation 3, we arrive at

$$-\frac{1}{\rho} p_{,ij} = \Omega_{ik} \Omega_{kj} + S_{ik} S_{kj}. \quad (4)$$

The eigenvalues of $S^2 + \Omega^2$ can be used to define vortical motion:

$$[S^2 + \Omega^2 - \lambda_i I] \mathbf{X}_i = 0, \quad (5)$$

where I is the identity matrix, and λ_i and \mathbf{X}_i represent the i^{th} eigenvalue and its corresponding eigenvector. A negative eigenvalue indicates that a “pressure minimum” has occurred in a plane defined by the corresponding eigenvector. Since a vortex core is not defined by a single plane, but rather by the intersection of two, we will look for the presence of two negative eigenvalues. Therefore, if we order the eigenvalues such that $\lambda_1 \leq \lambda_2 \leq \lambda_3$, the condition of $\lambda_2 < 0$ indicates the presence of a vortex.

It is advisable to non-dimensionalize λ_2 so that it remains well scaled for a variety of flow cases (similar to Section III.A), and we again use $\|S\|^2$ to define a threshold function

$$f_{threshold} = -\frac{\lambda_2}{\|S\|^2}. \quad (6)$$

The negative sign has been inserted so that a positive threshold ($f_{threshold} > 0$) value may be used to control refinement. Although it is a little less intuitive in this case, λ_2 and $\|S\|^2$ have similar scalings (and units) since both represent local velocity gradients.

III.C. Modified Δ

Chong²⁰ *et. al.* investigated how three-dimensional flow fields may be classified by the eigenvalues of the velocity gradient tensor. Using techniques of critical point theory, a local Taylor series expansion of the flow can be used to describe topological features of flow patterns. Similar to the eigenvalue problem of Equation 5, the eigenvalues of $\nabla\mathbf{u}$ reveal information about the local flow field and the solution to the corresponding problem for the 3D case is

$$\lambda^3 + P\lambda^2 + Q\lambda + R = 0, \quad (7)$$

where $P = -\text{trace}[\nabla\mathbf{u}]$ ($P = 0$ for an incompressible flow), $Q = (P^2 - \text{trace}[(\nabla\mathbf{u})^2])/2$, and $R = -\text{det}[\nabla\mathbf{u}]$. The solution to Equation 7 will either yield three real roots ($\lambda_{r1}, \lambda_{r2}, \lambda_{r3}$) or a single real root with a complex pair ($\lambda_r, \lambda_{cr} \pm \lambda_{ci}$). If the latter case occurs, local swirling motion exists. Moreover, the eigenvalues also represent the local kinematics (Figure 2).

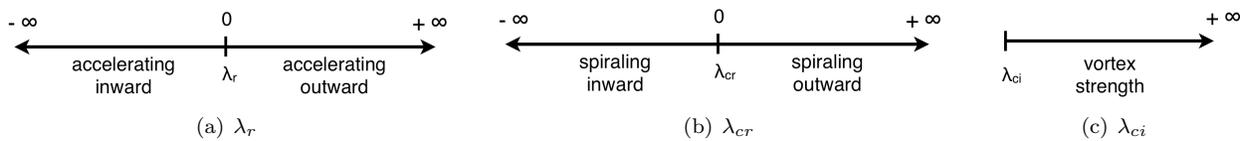


Figure 2: A demonstration of how eigenvalues of $\nabla\vec{u}$ effect particle motion.

The sign and magnitude of λ_r determine the acceleration and direction of the particle along a vector that is normal to the plane of vorticity; the sign and magnitude of λ_{cr} indicate whether or not the flow is swirling into or out of its center; and the magnitude of λ_{ci} indicates the swirling (vortical) strength. Since λ_{ci} is a measure of rotational strength, we expect a high correlation between it and the vorticity. However, to make it practical for our current purpose, this parameter is again non-dimensionalized by $\|S\|$, so the threshold

function is defined as:

$$f_{threshold} = \frac{\lambda_{ci}}{\|S\|}. \quad (8)$$

Note that λ_{ci} is non-dimensionalized by $\|S\|$ for proper dimensionality, and not by $\|S\|^2$ as was done for the previous cases.

III.D. Correlation between S and Ω

The turbulence community has expended considerable effort on classifying *coherent structures*, which are identifiable topologies found in turbulent flows that have similar observable characteristics. These structures may be roughly classified into two groups: those with vortex tube-like structures, and those with vortex sheet-like structures.²³ Although this classification is useful, it should be borne in mind that these groups often coexist, since a vortex tube is initially formed from the roll up of a vortex sheet.

The three previous methods have been designed to pinpoint vortex cores and here we examine a method that attempts to locate a vortex sheet. To identify regions of vortex sheets, we will leverage the fact that strain and rotation rates are correlated, and that both have large magnitudes in vortex sheets. Horiuti²¹ *et al.* suggested the eigenvalue problem of

$$[S\Omega - \Omega S - \lambda_i I]X_i = 0. \quad (9)$$

The term $(S\Omega - \Omega S)$ is also found in the anti-symmetric part of a simplified form of the Reynolds stress tensor.²⁴ Since this method looks for large positive eigenvalues contained in a vortex plane, the search is constrained to planes that are normal to the vorticity vector. Therefore, the eigenvalue associated with the eigenvector that is maximally aligned with the vorticity vector is ignored, and λ_+ (the second largest remaining eigenvalue) is used to control mesh refinement. This eigenvalue is also non-dimensionalized by $\|S\|^2$, but an offset is added since we are looking for regions where the vorticity strength is *greater* than the shear rate strength (as was the case for the non-dimensional Q method), so we define

$$f_{threshold} = \frac{\lambda_+}{\|S\|^2} - 1. \quad (10)$$

We reiterate that, for all the above methods, the threshold function is used to mark mesh cells for refinement if its value is greater than a pre-specified threshold value. Each threshold function has a zero offset so that all positive values indicate regions of swirl. In the following section, we investigate performance of the non-dimensional methods in the context of dynamically adapted solutions.

IV. Results

In related earlier work,¹⁰ a validation of the proposed non-dimensional schemes was performed by applying the methods to a series of static flow solutions. It was shown that the different non-dimensional methods were reliably robust and highly effective in identifying vortices under widely varying flow conditions, such as Mach number, vortex strength and size, vortex separation distance, and grid resolution. In the present work, we apply the methods to calculate dynamically adapted solutions within the Helios framework with the non-dimensional feature-detection methods being used to drive the off-body adaptive mesh refinement process.

The test cases used in this study are: a single incompressible Lamb vortex, two interacting incompressible Lamb vortices, a 3D ring vortex, and a NACA 0015 wing flow-field. A key objective is to verify that the tagging procedure is fully-automated and that it can be used without re-tuning the adaption criteria for a wide variety of problems. In particular, we stress the importance of using the proposed non-dimensionalized forms of the feature-detection methods and the specification of a *fixed* threshold value as a tagging criterion for each method. The traditional vorticity-magnitude method is also applied for purposes of comparison. In addition to solution accuracy, average execution run times are also evaluated.

IV.A. Single Incompressible Lamb Vortex

For this integrated case, an incompressible Lamb vortex given by

$$V_\theta = \frac{\Gamma}{2\pi} \left(\frac{1 - \exp^{-r^2/a^2}}{r} \right), \quad (11)$$

advects through the computational domain at a constant rate. Comparing the final solutions obtained by the adaptive methods to the exact solution represented on a fine grid will simultaneously reveal both dissipative and convective errors. A subsonic vortex with strength $\Gamma = 0.2\pi$ and radial size $a = 1.5$, is placed at the origin of a computational domain with $\mathbf{x} \in [-20, 20]$ and $\mathbf{y} \in [-20, 20]$. Three levels of refinement are used on an isotropic grid, where the mesh spacing is $h_{l_0} = 4/3$ (about 3 points across the vortex), $h_{l_1} = 2/3$ (about 5 points), and $h_{l_2} = 1/3$ (about 9 points). The vortex convects rightward ($+x$) at Mach = 0.1 and with periodic boundary conditions for x (and Neumann conditions in y), a single period is $t = 400$. With $\Delta t = 0.1$, we use 4,000 time-steps for the full simulation.

The regridding frequency is set *a priori* and is based on the speed of the vortex with respect to the finest grid spacing. We define the frequency (σ_R) as the ratio of the regridding time-interval to the simulation time-step size, i.e., $\sigma_R = \Delta t_G / \Delta t$. The re-gridding time-interval is then calculated by applying the following

requirement:

$$CFL_G = \frac{V_{ref} \Delta t_G}{h_{fine}} \approx 1. \quad (12)$$

Note that the parameter, CFL_G , is simply a CFL number based on the regridding time-interval. By setting this to a value of unity, we can ensure that the features of interest are properly tracked within the grid domain. For the current problem, when a CFL_G of unity is used, the regridding frequency $\sigma_R = 30$; in other words, regridding is performed every 30 time-steps.

Although the vortex is clearly identified when directly applying the non-dimensionalized methods, with $t_{val} = 1.0$, small spurious vortical motion, often close to the wall boundaries, is identified and tagged. Figure 3a is a snapshot of the grid and solution at $t = 960$ and illustrates what happens when the non-dimensional Q method is applied. (Note: similar behavior occurs for the other methods.) The contour threshold for Q

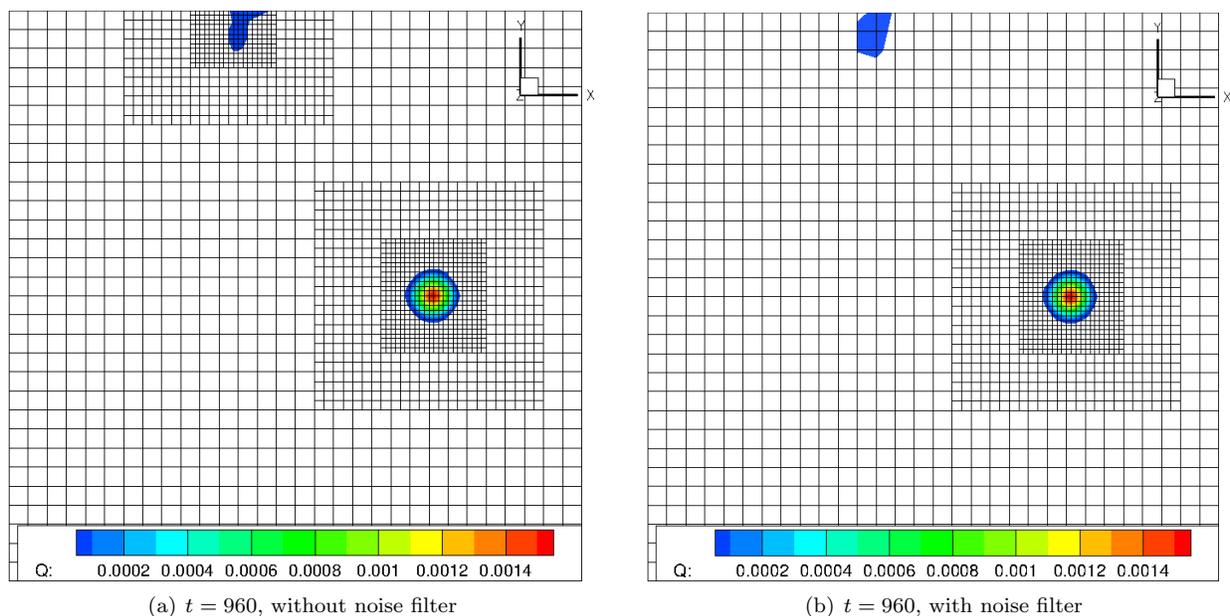


Figure 3: Non-dimensional Q method ($t_{val} = 1.0$) applied to advecting vortex case without (left) and with (right) a noise filter, where the contour levels of (dimensional) Q are shown.

is set to 1×10^{-9} to make this spurious vortex region noticeable, which is considerably smaller than the maximum value of 1×10^{-3} . Although small in magnitude, the local non-dimensional value of $\|S\|$ is near the same order of magnitude, which causes the method to tag this region.

These spurious regions cause tagging of non-critical features, which generates associated regions of grid refinement (Figure 3a). As a remedy in Figure 3b, a basic noise filter is implemented to avoid tagging a cell if the dimensional value, e.g., dimensional Q for the non-dimensional Q method, drops below a lower threshold value. To keep this lower bound (\bar{S}_{noise}) properly scaled with the problem, it is set to a percentage of the

global maximum dimensional value. Therefore, the cell is tagged if $f_{threshold} > t_{val}$ and the dimensional value exceeds \bar{S}_{noise} . For example, the non-dimensional Q method skips a cell if $Q_{local} < \bar{S}_{noise}$, where $\bar{S}_{noise} = \frac{\kappa}{100} Q_{max}$. Adding this noise filter does not cause degradation of the solution since the noise threshold values are kept low ($\kappa \approx 0.25$). Note that this does not significantly impact the automation of the method, and, therefore, a similar filter is applied to all the integrated test cases presented here.

To provide a comparison, the non-dimensional Δ , λ_2 , Q , and S - Ω correlation methods using a t_{val} of unity, the dimensional vorticity based method with $t_{val} = 8.67 \times 10^{-3}$, and a uniformly refined case are presented along with the exact solution after a full period ($t = 400$). The velocity magnitude is extracted along a line defined by $x = 0$, and Figure 4 demonstrates that all methods properly capture the convective speed of the traveling vortex and each has dissipated its strength by about 1%. We emphasize that the non-dimensional methods maintain generality by using a single threshold value that performs equally well for a wide range of vortex-based flows, in contrast to the vorticity threshold which is highly problem dependent. Moreover, all non-dimensionalized schemes show performance that is on par with uniform refinement, but they do so with about 10% of the total number of grid points, thereby reducing simulation time by almost an order of magnitude. Thus, the feature detection methodology promises to significantly reduce computational costs; a trend that is further amplified in the subsequent cases that use larger grid sizes.

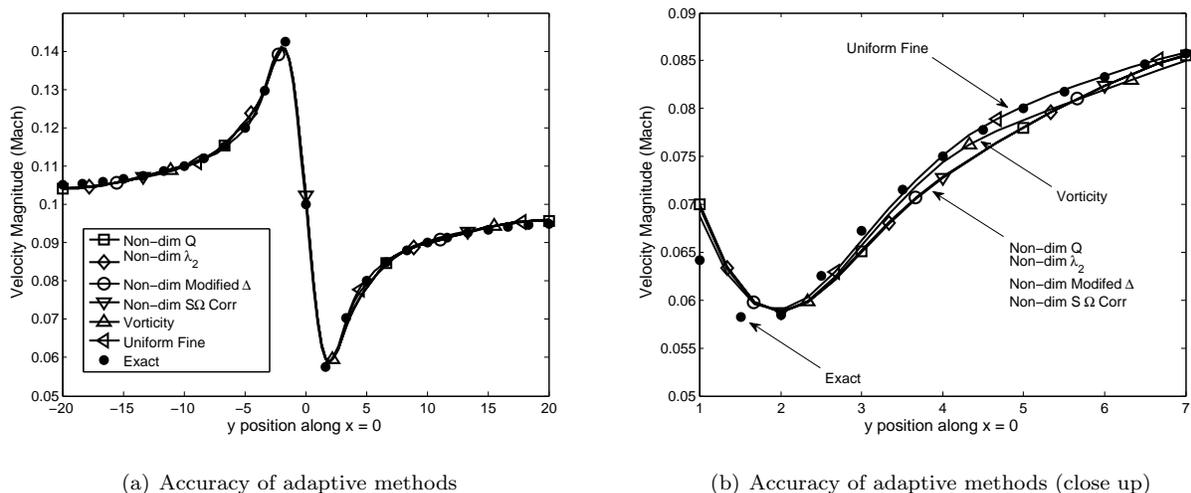


Figure 4: Comparison of methods for the integrated advecting vortex case after a single period. $t_{val} = 1.0$ for non-dimensional methods and $t_{val} = 8.67 \times 10^{-3}$ for the vorticity method.

IV.B. Two Incompressible Lamb Vortices

Rather than convecting the vortices along an axis, this test case takes advantage of a unique property of a co-rotating vortex pair. Meunier^{25,26} *et al.* have discovered that the rotational speed of a basic vortex system may be approximated by a simple relationship. When two vortices are separated by a distance that

is large compared to their core radius, the macro-scale inviscid dynamics of the system may be properly modeled by a point vortex approach. When the vortices are of equal strength Γ and separated by a distance of b_0 , they will rotate around each other at a constant angular speed of $\Omega = \Gamma/(\pi b_0^2)$. This vortex system behaves as if all vorticity is located at the vortex centroid (akin to a ‘center of mass’). For this case, two co-rotating vortices are placed at $\mathbf{x}_1 = \{5.513, 0\}$ and $\mathbf{x}_2 = \{-5.513, 0\}$, with strengths of $\Gamma = 0.2$ and core radii of 1.5 (yielding a quarter period of $t = 3,000$), and are placed in a computational domain with $\mathbf{x} \in [-25, 25]$ and $\mathbf{y} \in [-25, 25]$. Three levels of refinement are used on an isotropic grid, with $h_{l_0} = 1.0$, $h_{l_1} = 0.5$, and $h_{l_2} = 0.25$. Neumann boundary conditions are applied. A quarter period rotation requires 30,000 time-steps at $\Delta t = 0.1$, and re-gridding was done every 600 steps by enforcing a CFL_G of unity.

The non-dimensional methods are applied with t_{val} of unity while the dimensional vorticity method uses $t_{val} = 2.79 \times 10^{-3}$, which is again problem dependent. As an example, the λ_2 method (using $t_{val} = 1$) is applied, and Figure 5 shows the counter-clockwise rotating pair at $t = 13,800$ and $t = 30,000$. Although

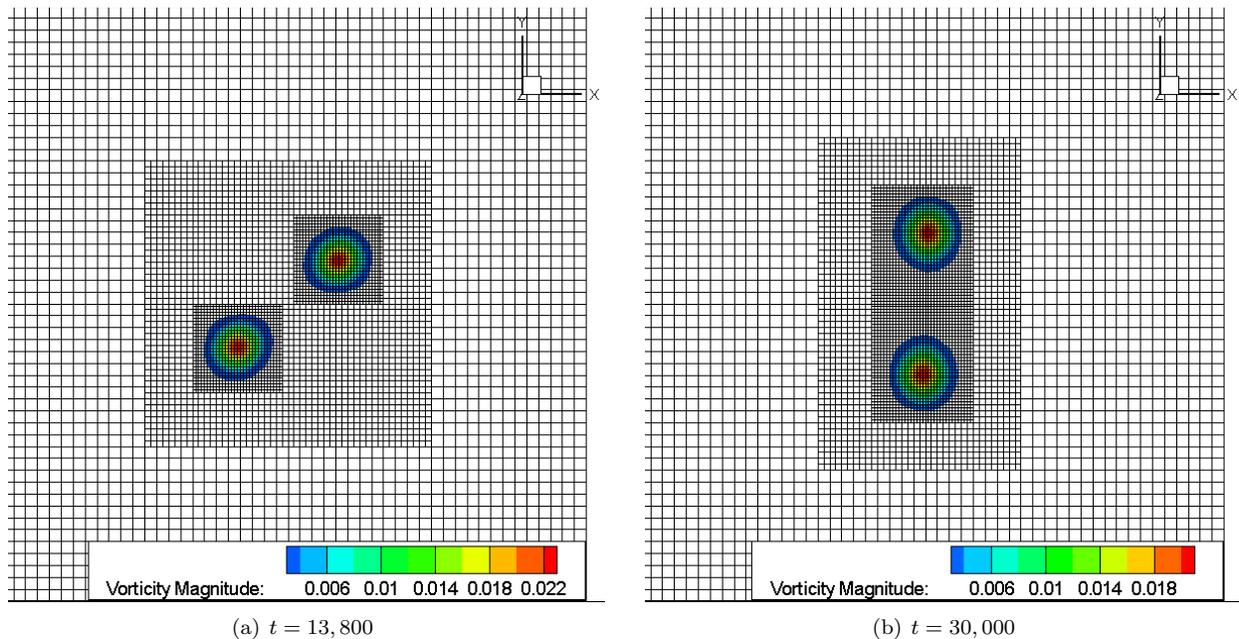


Figure 5: Non-dimensional λ_2 method ($t_{val} = 1.0$) applied to co-rotating vortex case visualized at approximately an eighth (left) and quarter (right) of a full period.

an exact solution only exists for a point vortex model, Le Dizès *et al.*²⁷ has shown that, for cases similar to the one considered here, where the vortex core sizes and separation distances are of appropriate lengths, the vortex system may be modeled as two point vortices. Having verified the existence of an analytic solution, the velocity magnitudes of the adaptive methods and the exact solution are compared along a line where $x = 0$ (Figure 6a). The flow-field has rotational symmetry, so we consider one of the two vortices with points extracted from the origin to the far-field. Note that, for the adaptive cases, the rate of rotation is marginally slower which causes the vortex cores to not precisely overlap with the 90° and 270°

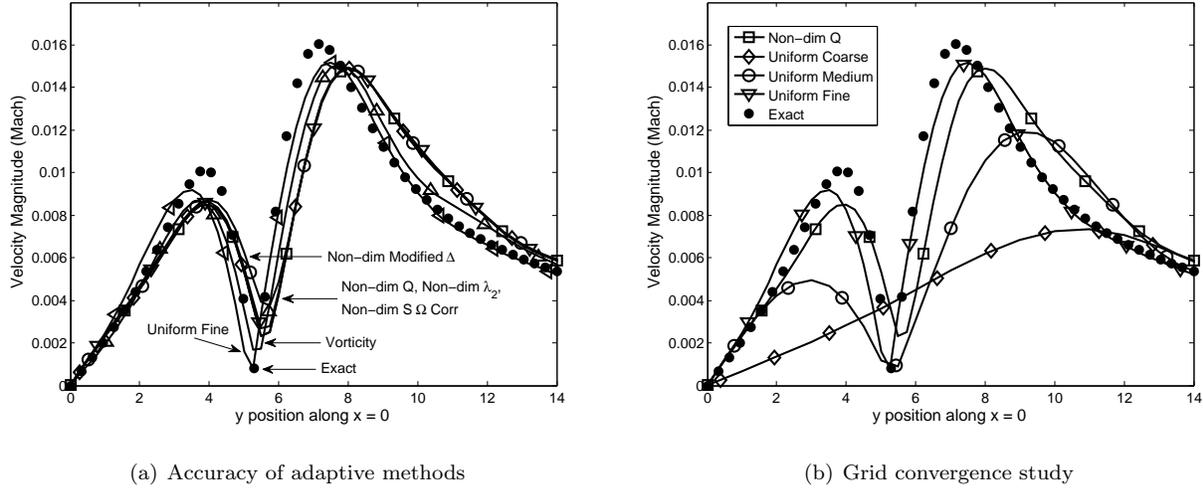


Figure 6: The co-rotating vortex case after a quarter period. $t_{val} = 1.0$ for non-dimensional methods and $t_{val} = 2.79 \times 10^{-3}$ for the vorticity method.

radial positions. This has the adverse effect of aliasing the slight convective error as a moderate diffusive error, since our 2D measurement does not accurately “sample” the maximum and minimum core velocities. This effect is lessened for the dimensional vorticity magnitude method where the finely tuned case-specific threshold value uses approximately 30% more points than the automated non-dimensional adaptive cases. Nevertheless, the non-dimensionalized feature detection algorithms yield similar performance to the uniform fine case. Although all the methods perform equally well, the non-dimensional Q , λ_2 , and $S-\Omega$ methods are slightly better at capturing the vortex minimum than the Modified- Δ method.

To provide a grid study comparison, the non-dimensional Q method is plotted against the solution obtained from the three uniformly refined grids (Figure 6b). The 2D vortex is simulated in a 3D domain, so the uniformly refined grids will contain approximately eight times as many nodes as their coarse counterpart, so each will take approximately eight times longer to run. Although the non-dimensional Q method is used as a reference, all adaptive methods outperform the uniform coarse and uniform medium cases and come close to approximating the uniform fine solution. More importantly, Table 1 demonstrates that the adaptive methods are approximately ten times faster than the comparable uniform-fine grid computation.

Table 1: Comparison of average run time per time-step and number of grid points for the co-rotating vortex case on uniform and adaptive grids. Adaptive grid size listed is calculated using the mesh at $t = t_{final}$.

	n	sec/time-step
Coarse Uniform	2.43×10^4	0.115
Medium Uniform	1.77×10^5	0.760
Fine Uniform	1.37×10^6	5.51
Non-dim Q (adaptive)	1.36×10^5	0.538

IV.C. Three-dimensional Ring Vortex

The potential flow solution for the ring vortex was calculated analytically by Lamb.²⁸ The entire ring moves at a constant rate along the z -axis and for a Rankine vortex, w_{ring} is given by

$$w_{ring} = \frac{\Gamma}{2\pi d} \left(\ln \left(\frac{4d}{a} \right) - \frac{1}{4} \right), \quad (13)$$

where d is the ring diameter and like before, a is the core radius and Γ represents circulation strength. For this case, the ring vortex was placed at the origin, with $a = 1.0$, $d = 6.0$, and $\Gamma = 0.5$. We apply Neumann boundary conditions to a domain with $\mathbf{x} \in [-20, 20]$, $\mathbf{y} \in [-20, 20]$, and $\mathbf{z} \in [-20, 40]$. Four levels of refinement are used on an isotropic grid with, $h_{l_0} = 2$, $h_{l_1} = 1$, $h_{l_2} = 0.5$, and $h_{l_3} = 0.25$ (Figure 7).

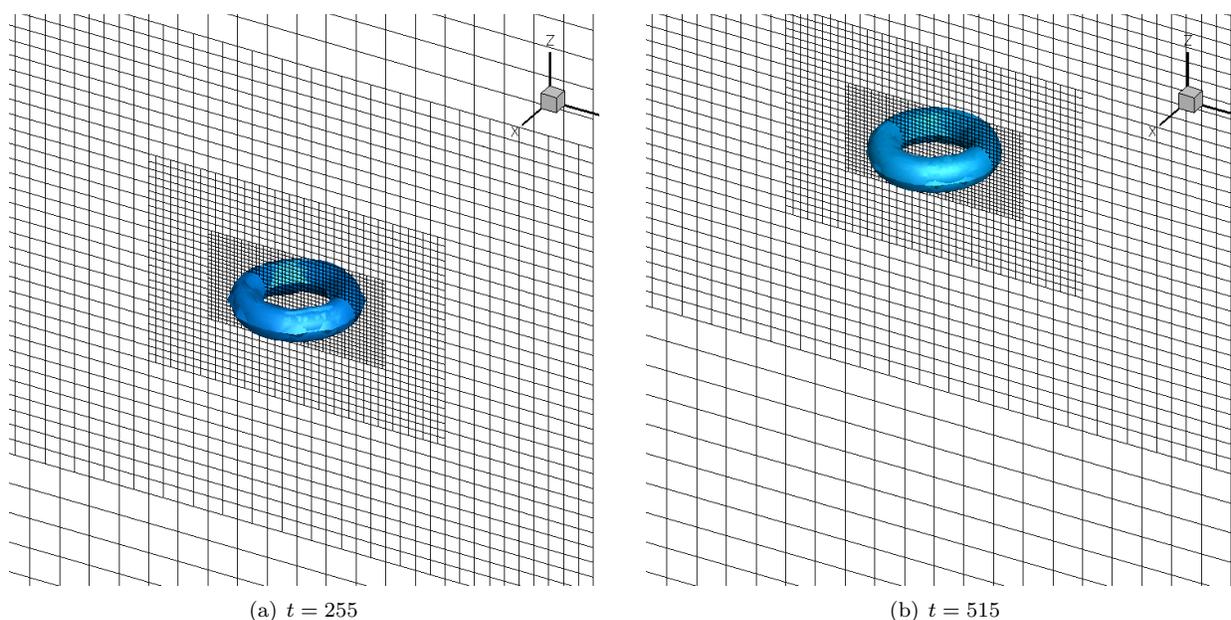
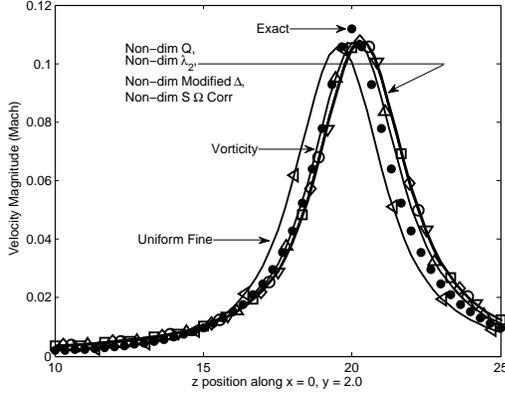


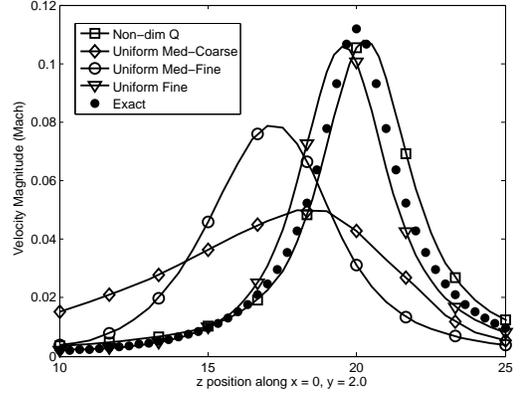
Figure 7: Modified- Δ method ($t_{val} = 1.0$) applied to 3D ring vortex case after advancing to $z = 9.9$ (left) and $z = 20$ (right).

At $t = 515$ the ring travels to $z = 20$, over the course of 5150 time-steps ($\Delta t = 0.1$) and a re-grid operation is carried out every 50 time-steps ($CFL_G \approx 1$). Each non-dimensional method with a t_{val} of unity, the dimensional vorticity with a t_{val} of 5.85×10^{-2} , and a uniformly fine case are compared against the exact solution. Figure 8a compares the velocity magnitudes at the end of the simulation along a line, normal to the x - y plane, that passes through the ring and is defined by $x = 0.0$ and $y = 2.0$.

Perhaps surprising, all adaptive grid methods outperform the uniform fine case (Figure 8b). It is somewhat fortuitous that the coarser grid regions favorably influenced the convection and dissipation of the the traveling ring vortex, and this will be further examined in future work. As before, the uniformly refined grids contain approximately eight times as many nodes as their coarse counterpart so each will take approx-



(a) Accuracy of adaptive methods



(b) Grid convergence study

Figure 8: Comparison of methods for the integrated 3D ring vortex case at $t = 515$. $t_{val} = 1.0$ for non-dimensional methods and $t_{val} = 0.0585$ for the vorticity method.

imately eight times longer to run (Table 2). All adaptive methods (non-dimensional Q used as a reference) demonstrate comparable accuracy to the uniform fine case, but they do so about 40 times faster. Note that this improvement is remarkable considering that more than 100 re-grid operations were performed during the simulation, which confirms the relatively low computational cost associated with the re-grid operations (compared to the expense of the flow solution). Although detailed temporal accuracy studies will be saved for future work, this result is again very promising.

Table 2: Comparison of average run time per time-step and number of grid points for the ring vortex case on uniform and adaptive grids. Adaptive grid size listed is calculated using the mesh at $t = t_{\text{final}}$.

	n	sec/time-step
Med-Coarse Uniform	1.07×10^5	0.461
Med-Fine Uniform	8.34×10^5	3.51
Fine Uniform	6.59×10^6	27.3
Non-dim Q (adaptive)	1.87×10^5	0.649

Furthermore, essential to developing an automated approach, the non-dimensional methods accurately track the vortex with a fixed threshold value of unity, while the value set for the method based on the vorticity-magnitude has to be modified on a case-by-case basis. The need to manually tune the traditional schemes will only become even more problem-dependent for practical applications such as the following NACA wing case.

IV.D. NACA 0015 Flow-Field

Having demonstrated the performance of the methods for a variety of theoretical cases, we now consider the steady flow around a full-span NACA 0015 square wing at a Mach number of 0.1235 with a 12° angle of

attack at a Reynolds number of 1.5×10^6 . In the farfield, Dirichlet boundary conditions are imposed with a uniform freestream. The experimental test was originally performed by McAlister²⁹ *et al.*, and our solution is obtained with the computational approach outlined in Section II. The hybrid unstructured-Cartesian grid system at the start of the simulation (before any adaptive refinement is applied) is illustrated in Figure 9. The unstructured grid encompasses the wing section, while the off-body Cartesian grid system is comprised

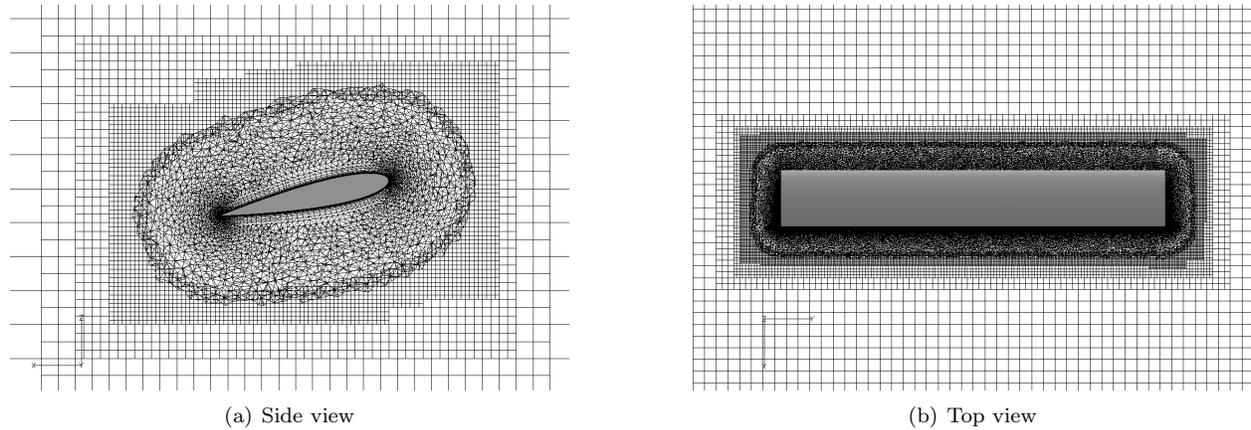


Figure 9: The NACA 0015 near-body unstructured/off-body Cartesian overset grid system.

of four mesh levels where $h_{l_0} = 0.2$, $h_{l_1} = 0.1$, $h_{l_2} = 0.05$, and $h_{l_3} = 0.025$. All grid units are non-dimensionalized by the chord length. The case was executed as a steady flow problem and regridding was performed every 250 iterations. In general, convergence was obtained after about 75 adaption steps.

In the non-adaptive case, notice that the trailing tip vortices are almost completely dissipated as they enter into the coarse grid region (Figure 10). Near the wing, the vorticity is relatively high, but at three

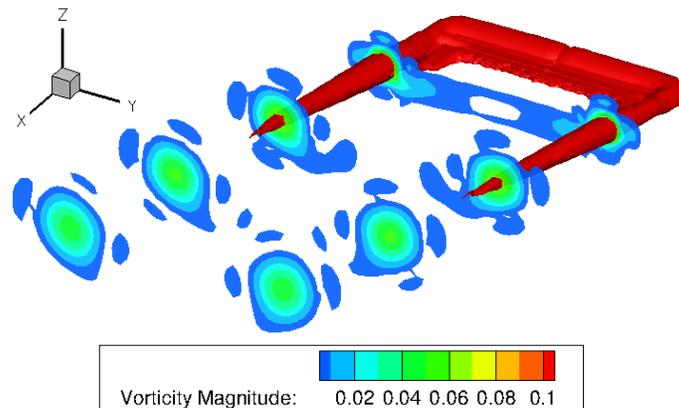


Figure 10: Vorticity isosurface at $\|\vec{\omega}\| = 0.1$ and four cutting y - z planes with vorticity contours for unadapted mesh.

chord lengths downstream, the flow only retains 3% of this value. This clearly emphasizes the need for localized grid refinement that preserves the swirling motion far downstream of the wing.

Figure 11 shows adapted results from different feature detection techniques. We reiterate that the mesh

adaption is applied to the off-body Cartesian domain, where cells are marked if $f_{threshold} > t_{val}$. Results are shown both for feature detection based on traditional methods such as the vorticity magnitude (11a) as well as two of the proposed non-dimensional methods (11b).

As in the previous theoretical examples, the determination of an appropriate threshold value remains a trial-and-error process when the vorticity magnitude is used. Figure 11a illustrates some of the variability and precision required by the dimensional-vorticity method by showcasing the corresponding grid refinement when two different threshold values are applied. For the higher threshold value ($t_{val} = 0.75$), the final adaptive grid system extends only a short distance downstream of the wing trailing edge because the trailing vortices are not adequately tagged. Only when the threshold is manually tuned to 0.25, do we obtain complete tagging of the vortex through the whole domain.

On the other hand, as seen in Figure 11b, the non-dimensional methods correctly tag the vortex without any additional effort. As in the theoretical cases, we apply a fixed threshold value of unity and both the non-dimensional Q and the Modified- Δ methods identify the entire trailing vortex system, and the AMR procedure correctly furnishes the proper level of refinement in these vortex-dominated regions. Note that the non-dimensional λ_2 and $S-\Omega$ correlation methods were not studied in this case, but we expect similar performance with these methods as well.

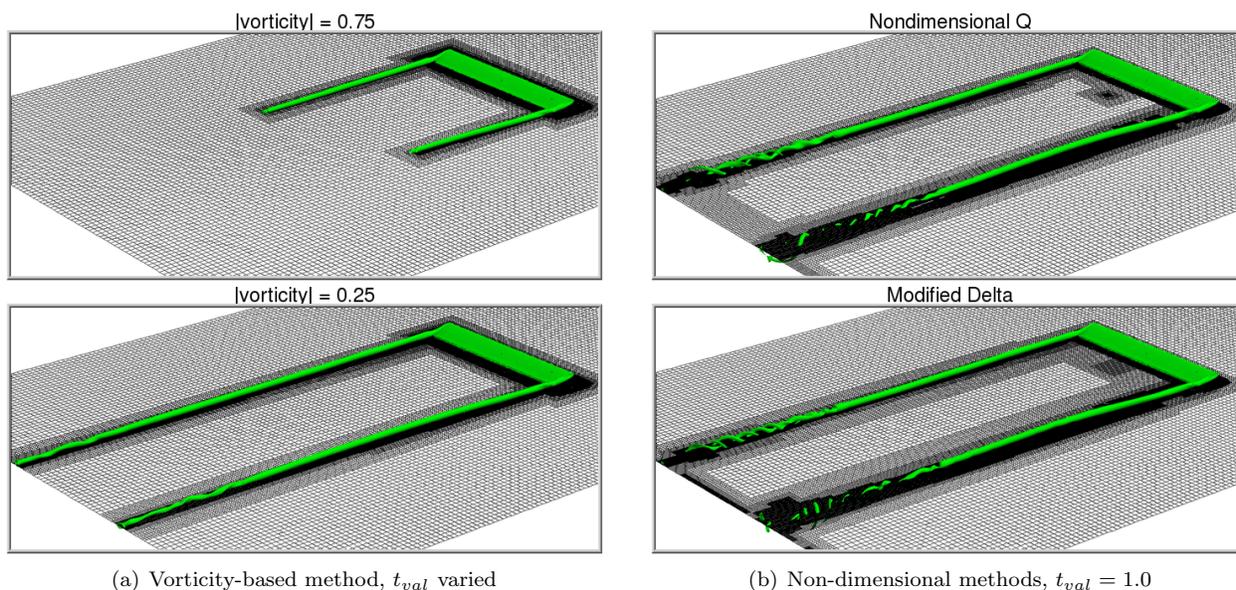


Figure 11: Grid system with corresponding solution at convergence for various adaptive methods for NACA 0015. Isosurfaces of trailing vortex motion represent vorticity magnitude.

Another outcome of the non-dimensional approach is that they tag a slightly broader region of the vortex when compared with the “optimal” vorticity magnitude method ($t_{val} = 0.25$). This further appears to allow for a better capture of the downstream vortex, as evidenced by the pronounced vortex instability at about twelve chord lengths downstream of the trailing edge. While further studies are needed to verify the accuracy

of these predictions, it is noteworthy that proposed non-dimensional methods appear to be more sensitive to the flowfield gradients than the vorticity magnitude method.

The enhanced sensitivity of the non-dimensional methods is also evident when looking at the relative progress of the grid adaption procedure during the course of the simulations. Figure 12 shows the results during the early stages of the convergence process for the properly tuned vorticity-based method ($t_{val} = 0.25$) and the non-dimensional Q method. It is noteworthy that the vortex development in the two instances are approximately the same, but the adaptive meshes produced by the two methods are dramatically different. Specifically, the vorticity method only tags cells with vorticity values of 0.25 or larger, and, therefore, the refinement is confined to the immediate vicinity of the vorticity iso-surface shown. On the other hand, the non-dimensional- Q method picks up more subtle levels of rotational flow and has already tagged the entire vortex path downstream of the trailing edge. This result also points to potential advantages of the non-dimensional techniques in terms of the rate of convergence for steady-state problems.

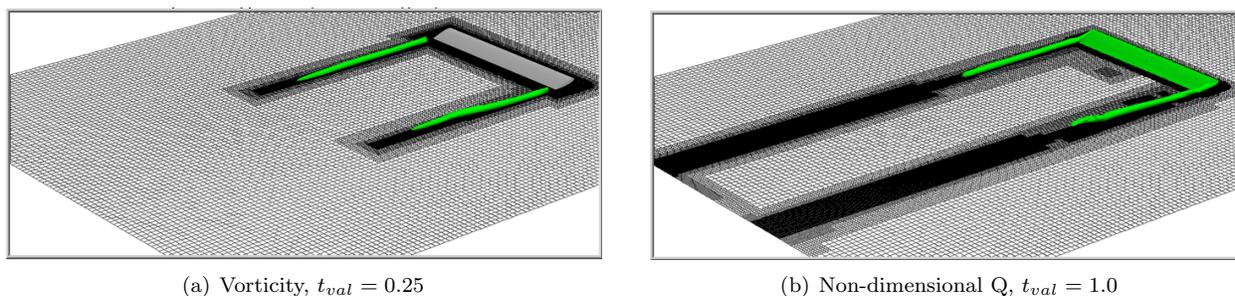


Figure 12: Unconverged grid system with corresponding solution for a dimensional (vorticity) and a non-dimensional (Q) method.

Another advantage obtained by considering the early stage of the solution in Figure 12 is that, for the non-dimensional methods, the size of the region tagged is directly proportional to the resolution of the vortex feature. Thus, vortices which are well resolved by the grid are tagged in a compact fashion, while vortices that are highly dissipated in coarse grid regions are tagged over a commensurately larger grid region. This suggests that the methods are automatically capable of readjusting the tagged region to properly capture the entire vortex. Indeed, this property may well be of even greater significance for unsteady problems, where certain unsteady features may be lost forever if they are not allowed to sufficiently develop into larger scale phenomenon when they initially occur in unrefined grid regions.

To quantify how effective the different methods are at preserving vorticity, the z -component of velocity (normalized by the freestream) is extracted from one of the vortex cores for cases with and without adaption, and Figure 13 compares these results with experimental data²⁹ at one and six chord lengths downstream from the trailing edge. The adaptive methods perform much better than the unrefined case and the non-dimensional methods fare as well as the optimally-tuned vorticity-magnitude method. However, even in the

adaptive cases, the maximum strength is lower than the reported experimental result.

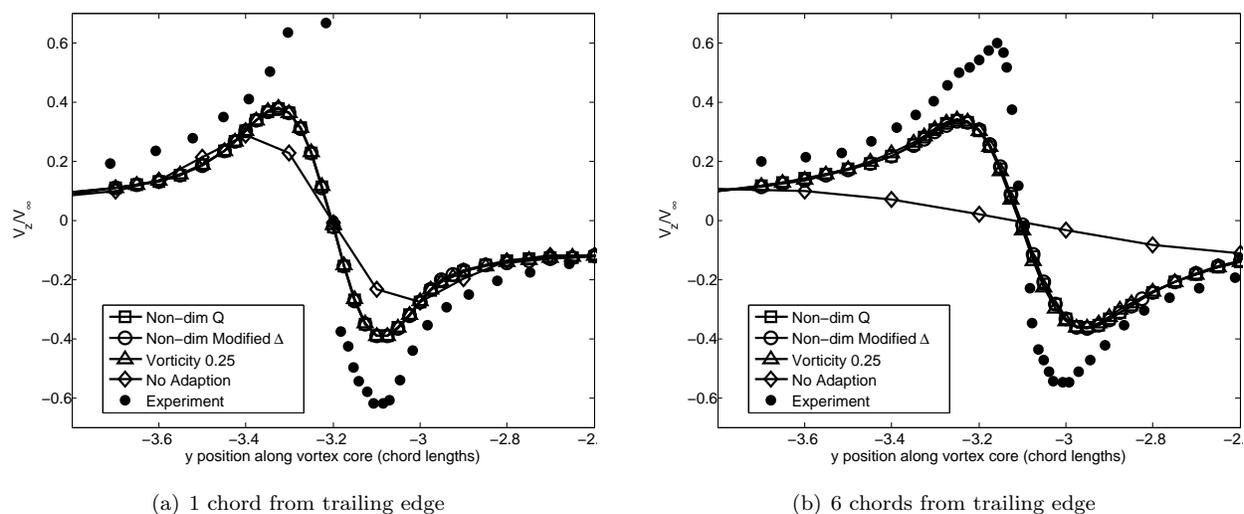
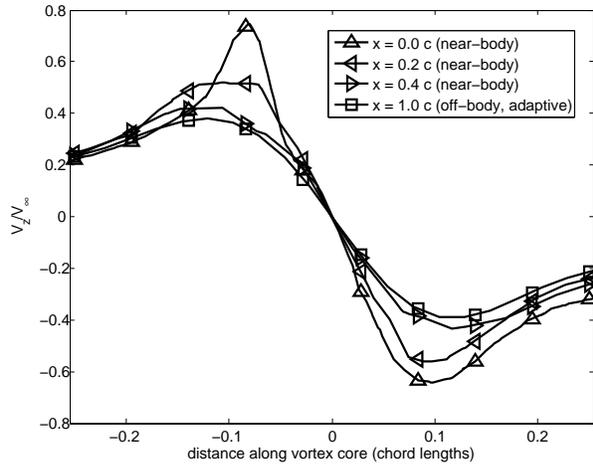


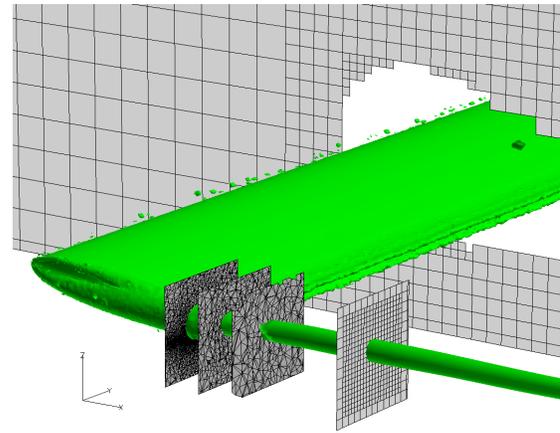
Figure 13: Normalized z -velocity for solutions using adaptive mesh refinement and no refinement compared against experimental results.

To help determine the root cause of the excessive vorticity dissipation, even with off-body adaption, we turn our attention to the vortex solutions within the near-body unstructured grid system. Figure 14 shows the z -component of the velocity at several downstream sections that are much closer to the trailing edge, where the vortex is still contained within the unstructured near-body mesh. Results are shown immediately downstream of the trailing edge, and at locations that are 0.2 and 0.4 times the chord length. In addition, we also include the results at one chord downstream at which the vortex is located in the off-body mesh. Comparing velocities found at the trailing edge and at 0.4 chord lengths downstream, we observe that the vortical strength is reduced by almost 50%, which indicates that much of the vortex dissipation occurs due to lack of adequate grid resolution in the near-body region. In fact, comparing the results at 0.4c (vortex in the near-body mesh) and 1.0c (vortex in the off-body mesh) only shows minimal vortex dissipation. Thus, it is clear that the adaptive off-body method is, in fact, performing very well at preserving the vortex strength once it is transferred to the off-body mesh. Future investigations will focus on increasing the near-body unstructured grid resolution to minimize the initial rapid vortex dissipation.

Figure 15 shows additional comparisons of the different adaptive methods as well as the base non-adapted scheme, focusing now on the peak normalized z -velocity component as a function of the downstream distance. The results indicate that the non-dimensional adaptive schemes, as well as the optimally-tuned vorticity method, perform very well in terms of preserving the velocity magnitude in the off-body domain. The small discrepancy of the Modified- Δ method at the twelve-chord location may be attributed to the unsteady vortex breakdown that was discussed earlier. In contrast, the poorly-tuned vorticity method decays very strongly



(a) Spanwise V_z/V_∞ at various locations along the vortex core.



(b) Location of four velocity measurements made at $x = 0.0c, 0.2c, 0.4c, 1.0c$. Isosurface drawn at vorticity = 1.0.

Figure 14: Details of near-body grid and near-field off-body solutions, showing normalized z -velocities extracted along a line passing through the vortex core. Flow is sampled at three near-body grid locations ($x = 0.0c, 0.2c, 0.4c$) and an off-body grid location ($x = 1.0c$). The adapted Cartesian off-body grid was obtained by applying the non-dimensional Q method.

over the first ten chord lengths while the case without refinement loses the vortex almost immediately due to the poor grid resolution of the base grid. These results offer further confirmation that the off-body adaption performs very well in maintaining the vortex magnitude downstream of the trailing edge. Moreover, it is clear that the traditional vorticity-based method yields results that are highly dependent on the selection of the optimal threshold value, which makes it impractical to use for general situations.

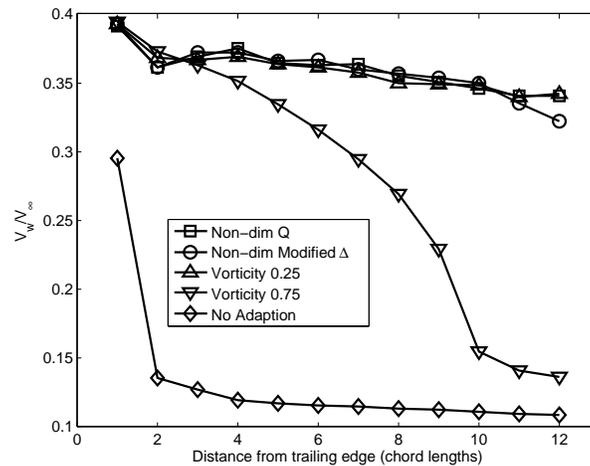


Figure 15: Normalized z -velocity for various solutions using adaptive mesh refinement and no refinement as a function of downstream distance from the trailing edge.

V. Concluding Remarks

This paper concerns feature detection techniques specifically designed for capturing vortex-dominated regions to drive Cartesian adaptive mesh refinement. Using traditional dimensional-based techniques to identify local regions for refinement, such as vorticity magnitude or the Q -criterion, can be problematic because the user must carefully tune the threshold that controls refinement. When the threshold is picked to be too large, important features are not identified and may dissipate quickly. When it is too small, minor features are identified for refinement and the domain is overpopulated with fine grid regions, thereby negating the benefits of localized AMR. Here, we introduce locally non-dimensionalized schemes for which a particular threshold can be commonly applied across a range of problems and flow regimes. A global noise filter is applied to the methods to remove spurious regions of refinement, and does so in a manner that does not significantly impact their efficiency or automation. Four particular schemes are evaluated: the non-dimensional Q , the non-dimensional λ_2 , the Modified- Δ , and the S - Ω correlation. The overall goal of this work is to evaluate these feature detection methods for vortex dominated flows and to demonstrate their use for general problems without need for user intervention.

Application of the developed methods are demonstrated for a hierarchy of tests, ranging from simple theoretical cases to complex practical flow-fields. The feature detection schemes are implemented as a module within the Helios framework, and the solution-based refinement strategy is used to drive the Cartesian off-body refinement. As highlighted by the theoretical cases (single, multiple, and ring vortices), the adaptive methods produce solutions that are comparable to uniformly fine meshes, but require only a small fraction of the computational cost. Extending this work to larger practical problems would potentially result in even more significant computational savings, due to larger areas of the grid where relatively coarser regions are acceptable.

In addition to these basic flow-fields, these methods were also applied to the flow over a NACA 0015 wing, where the trailing tip vortices are shown to be correctly identified over the full length of the computational domain. The vortex cores obtained from the non-dimensional methods without any specialized tuning were found to be nearly identical to those obtained using an optimally-tuned dimensional vorticity-based method. These results clearly underscore the general applicability, robustness, and efficiency of the proposed feature-detection methods.

Future work will focus on more practical applications of the feature detection methods within an integrated framework, particularly targeted towards rotorcraft aeromechanics computations, where unsteady effects will be addressed. Additionally, the consequences of improving the near-body grid resolution will be studied as well.

Acknowledgments

The first author and Prof. Jameson have been supported by the joint NASA-Stanford fellowship program (Grant #1125897-1-RAJJN). Material presented in this paper is a product of the CREATE-AV Element of the Computational Research and Engineering for Acquisition Tools and Environments (CREATE) Program sponsored by the U.S. Department of Defense HPC Modernization Program Office. Development was performed at the HPC Institute for Advanced Rotorcraft Modeling and Simulation (HI-ARMS) located at the US Army Aeroflightdynamics Directorate at Moffett Field, CA. The authors gratefully acknowledge the contributions to this work by Profs. Jay Sitaraman and Dimitri Mavriplis at the University of Wyoming and Dr. Thomas Pulliam at NASA Ames Research Center.

References

- ¹Wissink, A. M., Sitaraman, J., Sankaran, V., Pulliam, T., and Mavriplis, D., "A Multi-Code Python-Based Infrastructure for Overset CFD with Adaptive Cartesian Grids," 46th AIAA Aerosciences Conference, Vol. 927, January 2008, AIAA Paper 2008-927.
- ²Sitaraman, J., Katz, A., Jayaraman, B., Wissink, A. M., and Sankaran, V., "Evaluation of a Multi-Solver Paradigm for CFD using Unstructured and Structured Adaptive Cartesian Grids," 46th AIAA Aerosciences Conference, Jan. 2008, AIAA Paper 2008-660.
- ³Sitaraman, J., Wissink, A., Datta, A., Jayaraman, B., Potsdam, M., Sankaran, V., Mavriplis, D., Yang, Z., O'Brien, D., Saberi, H., Cheng, R., Hariharan, N., and Strawn, N., "Application of the Helios Computational Platform to Rotorcraft Flowfields," 48th AIAA Aerospace Sciences Meeting and Exhibit, Jan. 2010, AIAA Paper 2010-1230.
- ⁴Berger, M. and Olinger, J., "Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations," Journal of Computational Physics, Vol. 53, 1984.
- ⁵Berger, M. J. and Colella, P., "Local Adaptive Mesh Refinement for Shock Hydrodynamics," Journal of Computational Physics, Vol. 82, May 1989, pp. 64–84.
- ⁶Berger, M. and Jameson, A., "Automatic Adaptive Grid Refinement for the Euler Equations," AIAA, Vol. 23, No. 4, Apr. 1985, pp. 561–568, MAE Report No. 1633, October 1983, Princeton University, and NYU Report No. DOE/ER/03077-202, October 1983.
- ⁷Nemec, M., Aftosmis, M. J., and Wintzer, M., "Adjoint-Based Adaptive Mesh Refinement for Complex Geometries," 46th AIAA Aerosciences Conference, Jan. 2008, AIAA Paper 2008-725.
- ⁸Park, M. A., "Adjoint-Based, Three-Dimensional Error Prediction and Grid Adaptation," 32nd AIAA Fluid Dynamics Conference, 2002, AIAA Paper 2002-3286.
- ⁹Venditti, D. A. and Darmofal, D. L., "Grid Adaptation for Functional Outputs: Application to Two-Dimensional Inviscid Flows," Journal of Computational Physics, Vol. 176, 2002, pp. 40–69.
- ¹⁰Kamkar, S. J., Jameson, A. J., and Wissink, A. M., "Automated Grid Refinement Using Feature Detection," 47th AIAA Aerosciences Conference, Jan. 2009, AIAA Paper 2009-1496.
- ¹¹Mavriplis, D., "Unstructured Mesh Discretizations and Solvers for Computational Aerodynamics," 18th AIAA Computational Fluid Dynamics Conference, Jan. 2007, AIAA Paper 2007-3955.

¹²Pulliam, T. H., “Solution Methods in Computational Fluid Dynamics,” von Karman Institute for Fluid Mechanics Lecture Series, Numerical Techniques for Viscous Flow Computations in Turbomachinery, Jan. 1986, http://people.nas.nasa.gov/pulliam/mypapers/vki_notes/vki_notes.html.

¹³Hornung, R. D., Wissink, A. M., and Kohn, S. R., “Managing Complex Data and Geometry in Parallel Structured AMR Applications,” Engineering with Computers, Vol. 22, No. 3, 2006, pp. 181–195.

¹⁴Hornung, R. D. and Kohn, S. R., “Managing Application Complexity in the SAMRAI Object-Oriented Framework,” Concurrency and Computation: Practice and Experience, Vol. 14, No. 5, 2002, pp. 347–368.

¹⁵Wissink, A. M., Hornung, R. D., Kohn, S. R., S., S. S., and Elliott, N., “Large Scale Parallel Structured AMR calculations Using the SAMRAI Framework,” Supercomputing '01: Proceedings of the 2001 ACM/IEEE conference on Supercomputing (CDROM), 2001.

¹⁶Sitaraman, J., Floros, M., Wissink, A. M., and Potsdam, M., “Parallel Unsteady Overset Mesh Methodology for a Multi-Solver Paradigm with Adaptive Cartesian Grids,” 26th Applied Aerodynamics Conference, Jan. 2008, AIAA Paper 2008-7117.

¹⁷Wissink, A. M. and Shende, S., “Performance Evaluation of the Multi-Language Helios Rotorcraft Simulation Software,” Proceedings of the DoD HPC Users Group Conference, June 2008.

¹⁸Hunt, J. C. R., Wray, A. A., and Moin, P., “Eddies, Streams, and Convergence Zones in Turbulent Flows,” In its Studying Turbulence Using Numerical Simulation Databases, 2. Proceedings of the 1988 Summer Program p 193-208 (SEE N89-24538 18-34), Dec. 1988, pp. 193–208.

¹⁹Jeong, J. and Hussain, F., “On the Identification of a Vortex,” Journal of Fluid Mechanics, Vol. 285, 1995, pp. 69–94.

²⁰Chong, M. S., Perry, A. E., and Cantwell, B. J., “A General Classification of Three-Dimensional Flow Fields,” Physics of Fluids, Vol. 2, May 1990, pp. 765–777.

²¹Horiuti, K. and Takagi, Y., “Identification Method for Vortex Sheet Structures in Turbulent Flows,” Physics of Fluids, Vol. 17, No. 12, Dec. 2005.

²²Truesdell, C., The Kinematics of Vorticity, Indiana University Press, 1954.

²³Horiuti, K., “A Classification Method for Vortex Sheet and Tube Structures in Turbulent Flows,” Physics of Fluids, Vol. 13, Dec. 2001, pp. 3756–3774.

²⁴Horiuti, K., “Roles of Non-Aligned Eigenvectors of Strain-Rate and Subgrid-Scale Stress Tensors in Turbulence Generation,” Journal of Fluid Mechanics, Vol. 491, Sep. 2003, pp. 65–100.

²⁵Meunier, P., Ehrenstein, U., Leweke, T., and Rossi, M., “A Merging Criterion for Two-Dimensional Co-Rotating Vortices,” Physics of Fluids, Vol. 8, 2002, pp. 2757–2766.

²⁶Meunier, P., Le Dizès, S., and Leweke, T., “Physics of Vortex Merging,” Comptes Rendus Physique, Vol. 6, 2005, pp. 431–450.

²⁷Le Dizès, S. and Verga, A., “Viscous Interactions of Two Co-Rotating Vortices Before Merging,” Journal of Fluid Mechanics, Vol. 467, Feb 2002, pp. 389–410.

²⁸Lamb, H., Hydrodynamics, Dover Publications, 6th ed., 1993.

²⁹McAlister, K. W. and Takahasi, R. K., “NACA 0015 Wing Pressure and Trailing Vortex Measurements,” NASA Technical Paper, Vol. 3151, Nov. 1991.