Automated Off-Body Cartesian Mesh Adaption for Rotorcraft Simulations

Sean J. Kamkar *

Antony Jameson †

Stanford University, Palo Alto, CA 94305

Andrew M. Wissink ‡

Venkateswaran Sankaran §

Ames Research Center, Moffett Field, CA 94035

A new adaptive mesh refinement strategy is presented that couples feature-detection with local error-estimation. The goal is to guide refinement to key vortical features using feature detection, and to terminate refinement when a maximum acceptable error level has been reached. The feature detection scheme, which has been presented in previous related work, uses a special local normalization that allows it to properly identify regions of high vortical strength without tuning to a particular vorticity value. The newly introduced error estimation scheme applies a Richardson extrapolation-like procedure to detect local truncation error based on solutions from different grid levels. The error is then used the computed error to determine when to cut off further refinement. The paper presents a theoretical analysis of the scheme, applying it to computations of an isolated vortex and comparing to an exact solution. The scheme is implemented as part of the off-body Cartesian solver in the Helios code. Two practical cases are considered, resolution of the wake tip vortex from a NACA 0015 wing, and resolution of the wake structure of a quarter-scale V22 rotorcraft.

I. Introduction

Adaptive mesh refinement (AMR) is particularly useful for rotorcraft aerodynamics simulations where small scale features, such as tip vortices, that exist within a large computational domain can have a profound impact on vehicle performance, vibration, and noise. In spite of its advantages for this problem the use of AMR remains relatively uncommon in rotorcraft. Complexities like the inherently unsteady flow-field and the relative motion between the rotor and the fuselage make implementation of efficient adaptive schemes challenging. Specifically, while near-field blade-airload predictions have steadily improved over the past decade, obtaining accurate far-field wakes have proved to be much more computationally difficult.¹

Rotorcraft CFD is interested in the wake region for two primary reasons. First, computational studies have shown that accurate wake resolution leads to improvements in rotor performance predictions,² such as *figure of merit* used to measure the power efficiency of a rotor in hover. Second, the wake is largely characterized by the powerful vortex cores which are shed from the rotor blade tips. These regions can easily become entrained in hover and remain close to the fuselage in forward flight, leading to rotor-fuselage interactions that can be hazardous to the vehicle under certain flight maneuvers.³ Wake fidelity may be improved with polynomial (p) or spatial (h) refinement, and the present work leverages both by applying higher-order accurate discretizations coupled together with local mesh refinement.

Because the macro-scale dynamics of the wake can be sufficiently recovered through the proper resolution of its vortex cores, we focus on feature-detection techniques to guide the adaptive meshing strategy. Our previous work^{4,5} focused on the development of locally normalized feature detection algorithms. Normalized variations of the Q -criterion,⁶ pressure minima (λ_2),⁷ eigenvalues of the velocity gradient tensor (λ_{ci}),⁸ and the correlation between the symmetric (strain) and anti-symmetric (rotation) parts of the velocity gradient tensor⁹ were developed in this earlier work. The goal of these new feature detection schemes was

*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

to eliminate the extensive tuning required when adaptation is guided by traditional dimensional approaches such as vorticity magnitude or the Q -criterion. Optimal threshold values can vary significantly between different problems and therefore can require considerable tuning to properly adapt to wake features. The non-dimensional schemes were demonstrated to be successful in detecting vortical motion in the flowfield regardless of strength, size, and/or resolution. But their practical application to guide adaptation was limited by the fact that they identified all vortical features in the flowfield. Maximal grid refinement is not needed for every feature. Ideally, refinement should be added to features when the added resolution reduces solution error.

The objective of the current work is to develop a solution-based error estimator which may be coupled with the aforementioned non-dimensional feature detection schemes. The role of the error estimator is to control the amount of grid resolution to be applied, in effect turning off refinement of wake effects once the solution has reached a specified solution accuracy. The idea of using solution-based error outputs to control AMR is not new. Adjoint-based implementations^{10–12} have become standard practice guide AMR in steady problems. However, adjoint-based schemes introduce a number of complexities for unsteady problems, the primary one being that the adjoint problem must be fully solved backwards in time. This is non-trivial for large rotorcraft applications that may involve a quarter million time-steps or more^a, particularly because changes to the grid system introduced through adaptation would require multiple iterations of this procedure. Although attempts have been made to negate this cost through the use of quasi-steady approximations, Krakos and Darmofal¹³ have demonstrated that such techniques can lead to large perturbations in the reported error. Therefore, our approach seeks an error estimator based on local truncation error using Richardson Extrapolation in place of global adjoint-based error.

The remainder of the paper is organized as follows. A description of the adaptive overset CFD computational approach used for the present work is presented first in Section II. Next, Section III gives an overview of the feature detection scheme used to target candidate regions for mesh refinement. Section IV gives a theoretical analysis of the Richardson extrapolation (RE) error estimator together with spatial and temporal accuracy validation tests. Then Section V presents a strategy for coupling the feature-detection with the error estimator for automated and self-regulating AMR. Section VI presents results from application of the scheme to the NACA 0015 and quarter scale V-22 Tilt Rotor Aeroacoustic Model (TRAM) cases, comparing accuracy to available experimental data. Lastly, Section VII summarizes the major findings and discusses areas for additional work.

II. Computational Infrastructure

This work is developed as part of the CFD/CSD rotary-wing analysis package **Helios** (*Helicopter Overset Simulations*),¹⁴ being developed by the US Army’s High Performance Computing Institute for Advanced Rotorcraft Modeling and Simulations (HI-ARMS). To address the difficulties of accurate far-field wake resolution, **Helios** employs a dual-mesh overset flow solution approach comprised of an unstructured solver operating on the near-body prism/tetrahedral grid and a high-order adaptive Cartesian solver operating on the off-body wake domain.^{15,16} The two solvers are coupled using an overset domain connectivity algorithm that applies implicit hole cutting and adds fringe regions to allow for relative grid motion.¹⁴ The near-body grids can rotate and deform with the rotating grid system while the off-body grids remain stationary. See Figure 1. Further description of **Helios**’s main components are described elsewhere.

The purpose of this mixed near/off-body meshing strategy is to apply unstructured grids near the surface to resolve complex geometry and boundary layer effects, and block structured Cartesian off-body grids to resolve the far-field wake using a combination of high-order numerics and AMR. A similar paradigm is implemented in **OVERFLOW** and has achieved notable success.^{17,18} One key difference between **Helios** and **OVERFLOW** is the latter uses a near-body grid system composed of structured curvilinear grids. Also, **Helios**’s implementation couples disparate solvers whereas **OVERFLOW**’s near-body solver, domain connectivity, and off-body Cartesian solver are all within the same code.

The work in the present paper is largely motivated by the need to automate the off-body refinement process used in **Helios** by specifically targeting vortical features in the wake.

^aA calculation of 10 rotor revolutions using $\Psi = 0.1$ deg time-steps, which is common in calculations today, requires 360,000 time-steps

Figure 1: The Helios dual-mesh paradigm with unstructured near-body grids and the block-structured Cartesian grids in the off-body.

II.A. Off-body Cartesian Solver

The SAMARC solver¹⁹ combines the *Structured Adaptive Mesh Refinement Application Infrastructure* (SAMRAI) library^{20–22} developed by the Lawrence Livermore National Laboratory together with a block-structured ARC3DC Cartesian inviscid Euler solver with high-order discretizations developed at NASA Ames Research Center. SAMRAI manages the generation of meshes around flow features and geometries and handles the mesh partitioning, domain decomposition, and the *Message-Passing Interface* (MPI)-based parallel communication on distributed memory computer systems. The off-body domain is comprised of a multi-level Cartesian grid system stored as a union of rectangular blocks in the approach proposed by Berger and Collela.^{23,24} New-level construction occurs from coarsest to finest, and there is an agglomeration operation that groups regions of the same level together to create block-based refinement. Cells tagged for refinement will be organized in a manner to form rectangular blocks which contain finer grid levels above their coarser equivalents. See Figure 2). At each time step, the solution is first calculated on all blocks of a particular level.

Figure 2: Block-structured AMR grid containing a hierarchy of nested refinement levels.

Once complete, the finer level solution is injected into the next coarser level and the process continues on the coarser level. The use of the block-structured paradigm along with the algorithms that preserve computational efficiency in the SAMRAI library make adaptive solutions relatively cheap. Compared to a comparable unstructured grid, the Cartesian-based block structured grid storage requirement are considerably smaller. This small memory footprint allows for fast and efficient refinement, as demonstrated in other works^{25,26}

where SAMRAI-based refinement was applied to unsteady problems with refinement at every other time-step on systems with over 1,000 processors. Furthermore, the cost of the structured Cartesian grid solver is over an order of magnitude faster on a per-node compared to unstructured solvers. Additional details about the block structured solution in ARC3DC is provided elsewhere.⁵

III. Feature Detection

Several possible feature detection methods are summarized in our previous work.⁵ Here, we provide an overview of the method based on the Q -criterion⁶ which is used in the current study. Hunt⁶ originally developed the notion of Q , which is usually defined as, $Q = \frac{1}{2} (||\Omega||^2 - ||S||^2)$, where S and Ω respectively refer to the strain and rotational components of the velocity gradient tensor $(\nabla \mathbf{u})$. Although this definition measures the difference between the rotation and strain rate magnitudes, this value is still dependent upon the characteristic length and velocity scales inherent to the specific problem. To yield a suitable non-dimensional form, if normalized by $||S||^2$, a threshold function of

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

is obtained. Note that the threshold function is evaluated in each cell which is marked for refinement if the resulting functional value is greater than a pre-specified t_{val} , i.e., tagged if $f_{threshold} > t_{val}$. The normalization enforces automation by keeping appropriate threshold levels (t_{val}) problem independent. Irrotational flow occurs when $f_{threshold} \rightarrow -\frac{1}{2}$, solid body rotation happens when $f_{threshold} \rightarrow \infty$, and all positive values represent regions of swirl, i.e., regions where the rotational strength is larger than the shear rate strength. Although $f_{threshold}$ is unbounded when $||S|| \rightarrow 0$, it is acceptable since tagging will routinely occur when the function exceeds a finite threshold value (t_{val}). In the following section, we develop the error estimator that will be used to control the amount of resolution as found by Equation 1.

IV. Error Estimation

Feature detection schemes are very good at identifying regions for refinement, but they do not indicate when sufficient mesh resolution has been attained. In this section, we consider the use of the Richardson error (RE) as a termination criterion for the AMR procedure. We start with an examination of the *local truncation error*, or *LTE*. The *LTE* is expected to converge at rates governed by the order-of-accuracy of the applied spatial and temporal discretizations as the mesh is progressively refined. Furthermore, it represents the best attainable error convergence, and will therefore be used as a performance benchmark for the RE error estimator.

IV.A. Local Truncation Error

The general formulation is derived first, following which the *LTE* of an unsteady advecting vortex will be examined.

IV.A.1. LTE Formulation

For simplicity, the *LTE* is derived for the simple 1D linear wave equation,

$$\frac{\partial u}{\partial t} + \lambda \frac{\partial u}{\partial x} = 0, \quad (2)$$

where u is a general scalar that represents the exact solution, λ represents the wave-speed, and t and x represent the temporal and spatial dimensions, respectively. To solve Equation 2 numerically, it first must be discretized, so a discrete variable, w , is defined and initialized to u at $t = t_0$. For illustrative purposes, Equation 2 is temporally discretized with Forward Euler and spatially discretized with first-order upwinding. Now, the discrete difference equation at an arbitrary node j can be written as

$$w_j^{n+1} = w_j^n - \frac{\lambda \Delta t}{\Delta x} (w_{j+1}^n - w_j^n), \quad (3)$$

where Δt and Δx respectively represent the time-step and cell size. Now, the discrete residual operator, $\mathcal{R}(\cdot)$, is introduced and defined as the second term on the right hand side of Equation 3, i.e., $\mathcal{R}(w_j^n) = (w_{j+1}^n - w_j^n)$. The $\mathcal{R}(\cdot)$ operator contains the numerical flux balance for cell interface (or node) j , and drives the iterative solution process.

Since the *LTE* measures how well the discrete equation models the exact solution for a single time-step, both solutions are required to calculate it. Since w^{n+1} is guaranteed to satisfy the difference equation, it is likely that the exact solution, when inserted into $\mathcal{R}(\cdot)$, will not. The resulting non-zero term is the *LTE*, which for Equation 3 can be written as

$$LTE(x, t) = \frac{u(x, t + \Delta t) - u(x, t)}{\Delta t} + \frac{\lambda}{\Delta x} (u_{j+1} - u_j).$$

In turn, we note that this is equivalent to:

$$LTE = \frac{u^{n+1} - w^{n+1}}{\Delta t}.$$

Since the *LTE* is dependent upon space and time, it is usually expressed in terms of their delta quantities when discussing the numerical order-of-accuracy. For example, if the $LTE = \mathcal{O}(\Delta t^q) + \mathcal{O}(\Delta x^p)$, it is said that the numerical scheme is q^{th} - and p^{th} -order accurate in time and space, respectively. It can be shown that the above analysis is equally applicable to multi-stage Runge-Kutta schemes such as those used in the ARC3DC solver.²⁷ Such behavior is helpful when trying to understand how spatial and/or temporal refinement will work to reduce the error present in the solution and is examined in the following section.

IV.A.2. Validation

The *LTE* estimator is used to calculate the convergence behavior for an advecting Shu-Erlebacher-Hussaini²⁸ vortex. The solution is advanced for a single time-step from its initial state, which allows for a complete analysis of the spatial and temporal convergence as a function of the mesh size and time-step.

A vortex is placed at the center of a domain defined by $x \in [0, 32]$ and $y \in [0, 32]$, and the parametric study is conducted with $\Delta t = \{5, \frac{5}{2}, \frac{5}{4}, \dots, \frac{5}{512}\}$ and $\Delta x = \{2, 1, \frac{1}{2}, \dots, \frac{1}{32}\}$. For reference, the vortex convects through the domain at 0.1 Mach, which, given the size of the applied time-steps, translates the vortex at core widths ranging between 0.5 and 9.76×10^{-3} . Also, putting grid resolution into context, the vortex core radius is approximately 1.18, which indicates that the coarsest grid contains a single point in the vortex, while the most refined grid contains 75 points along the core diameter. Examining such a wide range of scales ensures that the asymptotic solution limit is attained for both Δt and Δx , so optimal error behavior is understood.

Figure 3a contains the spatial accuracy study, where the x -momentum error is plotted as a function of Δx . For large time-steps, the temporal error dominates the spatial error and the error stalls for constant Δt . However, when Δt is small enough, or when $\frac{\Delta t}{\Delta x}$ is approximately less than or equal to unity, fifth-order accuracy is clearly demonstrated. Note that this ratio represents the *CFL* number used in ARC3DC and is bounded by RK3's *CFL* limit of $\sqrt{3}$, which explains the optimal convergence near a *CFL* value of one.

The complementary temporal accuracy study is shown in Figure 3b. Although it showcases the same data as presented in Figure 3a, the lines now represent constant Δx to illustrate temporal convergence. Third-order accuracy is achieved in regions where $\frac{\Delta t}{\Delta x}$ is approximately equal to or greater than unity. Also, stalled convergence occurs when Δt is relatively large compared to Δx . In addition to validating the convergent behavior of the *LTE* for this vortex dominated flow, the study additionally verifies that the algorithms in ARC3DC are indeed third-order accurate in time and fifth-order accurate in space.

IV.B. Richardson Extrapolation

Although the *LTE* can express the discretization error its practical use is limited because the exact solution must be known. The RE method provides an effective substitute to the exact solution for practical problems of interest. In the previous *LTE* analysis it was demonstrated that the truncation error reduces at characteristic logarithmic rates as the mesh is refined. If the solution is in the asymptotic range the progression of finer meshes leads to an approximation of the exact solution. In the *LTE* analysis, the exact solution on a coarse mesh can be determined by applying a coarsening operation to the fine-mesh solution where

Figure 3: Convergence of the local truncation error (*LTE*) for the advecting vortex case for the L_2 norm of x -momentum error. The plots demonstrate that for $\frac{\Delta t}{\Delta x} \approx 1$, the scheme is fifth- and third-order accurate in space and time, respectively.

the nodal locations are coincident. This strategy represents the basic idea behind the RE error estimator. The resulting error is a *relative* error between the fine and coarse grid levels. A fine-mesh error can also be obtained by interpolating the coarse error

IV.B.1. Necessary Conditions

Assumptions for Richardson Extrapolation (RE) include 1) uniform and systematic refinement, 2) smooth and asymptotic solutions and 3) dominant discretization error. The work by Roy and Oberkampf²⁹ is recommended for a more comprehensive discussion of these assumptions.

1. The block-wise refinement used in **SAMARC** allows an assumption of uniform refinement. Enforcing systematic refinement in the Cartesian grid system is not an issue because cell aspect ratio, cell skewness, and cell-to-cell stretching all remain constant.
2. The condition of smooth and asymptotic solutions are generally valid for rotorcraft flowfields once they are sufficiently resolved. That is, it is unlikely that a coarse baseline mesh will contain a fully valid asymptotic solution of the wake but as a sequence of increasingly refined meshes are applied the solution will become asymptotic. Note that the error estimator is being used to *terminate* refinement, which is only likely to occur after the vortex cores are sufficiently resolved to the point they are in the asymptotic range.

The hierarchical block-structured mesh storage paradigm employed by **SAMARC** makes the RE estimator approach straightforward to implement since the Cartesian grid levels of different refinement are all retained in core memory. Hence, the necessary grid-management tasks associated with the RE algorithm are handled efficiently.

IV.B.2. Derivation

A Taylor series approximation can be used to express the exact solution in terms of the discrete solution and higher-order terms. For example, the discrete solution, w_h , is computed on a grid with spacing $h = \Delta x$,

$$u = w_h + \sum_{i=0}^{\infty} C_i h^{p+i}, \quad (4)$$

where w_h signifies a discrete solution on a grid where p refers to the applied spatial order-of-accuracy, and C_i represents a set of constants that are usually unknown. Applying Landau notation, the higher order terms

can be grouped together so that a single term is extracted from the infinite sum,

$$u = w_h + C_0 h^p + \mathcal{O}(h^{p+1}). \quad (5)$$

If the higher order term is dropped, Equation 5 represents a single linear equation. This equation can be used to create a system of linearly independent equations for each grid level available, which when combined, can be used to develop a better approximation of u . Equation 5 uses a grid size of h , so for demonstrative purposes, it is assumed that a finer solution is available on $h/2$,

$$u = w_{h/2} + C_0 \left(\frac{h}{2}\right)^p + \mathcal{O}\left(\left(\frac{h}{2}\right)^{p+1}\right). \quad (6)$$

Multiplying Equation 6 by 2^p and subtracting Equation 5, the equations combine to form a new estimate of u which relies on the data from the overlapping grids,

$$u = \frac{2^p w_{h/2} - w_h}{2^p - 1} + \mathcal{O}(h^{p+1}),$$

which improves the order by $\mathcal{O}(h)$ as compared to the original discrete representation (Eqn. 4).

In addition to improving solution accuracy, a similar technique can be used to compute an error estimate. If C_0 is solved for, the error on $w_{h/2}$ can be defined as the difference between the exact and computed solutions, i.e., $u - w_{h/2}$

$$\mathcal{E}_{h/2} = \frac{w_{h/2} - w_h}{2^p - 1} + \mathcal{O}(h^{p+1}). \quad (7)$$

Note that the denominator is a function of p (applied spatial order-of-accuracy) and is constant on all grid levels. As such, the error can be simply expressed as the difference between the coarse and fine solutions.

IV.B.3. Time-Accurate Considerations

For steady problems, Equation 7 can be directly applied to compute the error between two systematically refined grids. However, unsteadiness may influence both the spatial as well as the temporal errors.

It is assumed that the fine and coarse solutions are separated by n levels of refinement. Therefore, two solutions are available — a coarse solution on h and the finest level solution on $h/2^n$, where $n = \{1, 2, 3, \dots\}$. The wave equation is reconsidered (Eqn. 2), and if a discrete second-order accurate solution is constructed, the difference between the fine and coarse solutions is

$$\begin{aligned} w_f^{n+1} - w_c^{n+1} &= u_f^n - \lambda \Delta t \left. \frac{\partial u_f^n}{\partial x_f} \right|_d + \frac{\lambda^2 \Delta t^2}{2} \left. \frac{\partial^2 u_f^n}{\partial x_f^2} \right|_d \\ &\quad - u_c^n + \lambda \Delta t \left. \frac{\partial u_c^n}{\partial x_c} \right|_d - \frac{\lambda^2 \Delta t^2}{2} \left. \frac{\partial^2 u_c^n}{\partial x_c^2} \right|_d. \end{aligned} \quad (8)$$

Assume that the size of the time-step taken on the coarse and fine grids is identical and that the flow solution has gone through a single time-step from the initialization of u_f and u_c , so $u_f^n - u_c^n = 0$. SAMARC injects the finest solution available into its overlapping coarser level before the start of each time-step. Consequently, the error is computed at any point during the simulation represents the error generated during a single time-step. The fine-coarse difference is composed of the differences between the fine and coarse discrete derivatives,

$$w_f^{n+1} - w_c^{n+1} = -\lambda \Delta t \left(\left. \frac{\partial u_f^n}{\partial x_f} \right|_d - \left. \frac{\partial u_c^n}{\partial x_c} \right|_d \right) + \frac{\lambda^2 \Delta t^2}{2} \left(\left. \frac{\partial^2 u_f^n}{\partial x_f^2} \right|_d - \left. \frac{\partial^2 u_c^n}{\partial x_c^2} \right|_d \right). \quad (9)$$

For illustrative purposes, second-order accurate derivatives can be used so that the first bracketed term on the right hand side will resemble

$$\left. \frac{\partial u_f^n}{\partial x_f} \right|_{\Pi} - \left. \frac{\partial u_c^n}{\partial x_c} \right|_{\Pi} = \frac{\Delta x_f^2}{6} \frac{\partial^3 u_f^n}{\partial x_f^3} - \frac{\Delta x_c^2}{6} \frac{\partial^3 u_c^n}{\partial x_c^3} + \mathcal{O}(\Delta x_f^3) + \mathcal{O}(\Delta x_c^3).$$

Applying the assumption that $\Delta x_f = \Delta x_c / 2^n$, the above equation can be expressed in terms of Δx_c ,

$$\left. \frac{\partial u_f^n}{\partial x_f} \right|_{\Pi} - \left. \frac{\partial u_c^n}{\partial x_c} \right|_{\Pi} = \left(1 + \frac{1}{2^n} \right) \left[\frac{\Delta x_c^2}{6} \left(\frac{\partial^3 u_f^n}{\partial x_f^3} - \frac{\partial^3 u_c^n}{\partial x_c^3} \right) + \mathcal{O}(\Delta x_c^3) \right]. \quad (10)$$

When the flow is smooth, the higher-order derivatives will contain errors that are significantly smaller than Δx_c . For a general p^{th} -order accurate derivative, Equation 10 may be approximated as

$$\left. \frac{\partial u_f^n}{\partial x_f} \right|_{\text{P}} - \left. \frac{\partial u_c^n}{\partial x_c} \right|_{\text{P}} \approx \Delta x_c^p.$$

Now, the fine-coarse difference expressed by Equation 9 can be order-approximated by the following series,

$$w_f^{n+1} - w_c^{n+1} = \mathcal{O}(\Delta x_c^p) \sum_{k=1}^{\infty} \mathcal{O}(\Delta t^k). \quad (11)$$

Note that term on the right-hand side indicates that the spatial and temporal effects are still coupled.

Compare Equation 11 with the corresponding expression found by *LTE* analysis and it is found that the two uncoupled spatial and temporal terms no longer exist for the RE estimator. Instead, the Richardson estimate is completely comprised of coupled space-time terms, which will effectively remove the small- Δt stalled convergence which occurred for the *LTE*. Adding a $\frac{1}{\Delta t}$ factor to Equation 11 and assuming small non-linearities, we arrive at an expression for the error that is independent of the time-step applied,

$$\frac{w_f^{n+1} - w_c^{n+1}}{\Delta t} = \mathcal{O}(\Delta x_c^p) \sum_{k=0}^{\infty} \mathcal{O}(\Delta t^k), \quad (12)$$

where on the right-hand side, the small, yet important, summing variable switches from one to zero.

Either of the derived error expressions Eqns. 11 or 12 will recover proper p^{th} -order accuracy. The particular order expressed is a function of the *coarse* mesh spacing, indicating that an error computed between a coarse and a very-fine mesh will be similar in magnitude to the error computed by two meshes separated by a single level of refinement. Even though the error converges at a rate governed by the spacing on the coarse mesh it is still appropriate to use $w_f - w_c$ to estimate the error on the fine mesh in regions where points are co-located with coarser meshes. For those fine mesh points that lie between coarse mesh points interpolation may be used.

IV.B.4. Validation

Validation of the RE estimator uses the same advecting vortex problem from the earlier validation of the *LTE*. However, the number of refinement levels studied is one less because the extrapolation process can only generate $n - 1$ error estimates from n solutions. Figure 4(a) plots the RE error computed using Equation 12 comparing x -momentum solution error at different resolutions with the finest available ($\Delta x = \frac{1}{32}$). 5th-order spatial convergence is observed $\Delta x \rightarrow 0$ for all values of Δt , and the error decay rate does not stall, even for large Δt . As predicted, the spatial convergence remains generally unaffected by variations to time-step size. However, at the largest time-step of $\Delta t = 5$, the computed error for the coarsest grid ($\Delta x = 2$) appears to be slightly unstable. Note that the demonstrated convergent rates are in agreement with the *LTE* analysis presented earlier.

Figure 4(b) shows the convergence of the error measured between a level and the next finer level solution. The convergence looks essentially identical to the error between a level and the finest available (Fig. 4(a)). The spatial error behavior remains relatively asymptotic, even for the under-refined cases. The monotonically decreasing error in both cases indicates the inherent numerical stability associated with the RE method and its efficacy as a stable and robust control mechanism for solution adaption. Temporal accuracy studies presented elsewhere²⁷ show similar convergence.

One advantage of the RE estimator is that most other easily computed functions can be used to measure error and exhibit similar convergence rates as state variables ($\rho, \rho u, e$). Although it has been omitted, similar performance has been obtained when using Mach number, pressure, and velocity. For vortex-dominated flows, functionals that contain local maxima (or minima) near vortex centroids are preferred to target refinement at those regions. A pressure-driven approach has demonstrated to be more favorable than others at achieving this task, and the resulting local error behavior is provided next.

(a) Comparing to the finest available solution

(b) Comparing to the next-refined level

Figure 4: Spatial convergence of the Richardson estimator. Error in computed x -momentum for advecting vortex as compared to (a) *finest solution available* and ($\Delta x = \frac{1}{32}$) and (b) solution on *next finer level* for 5th-order spatial discretization.

IV.B.5. A Pressure-Based Error Estimator

Although the RE estimator demonstrates ideal error norm convergence, the prior convergence studies categorized global trends and therefore neglected local behavior. Here, a pressure-based estimator is examined by considering two analogous contour plot of the same vortex — one of pressure and the other of its associated error. The current test is performed on a vortex with about ten points across its core, and travels to the right at distance of $\Delta x = 0.1$ (a tenth of its width) at a speed of Mach 0.1. The solution and the associated error are presented in Figure 5. Examining the error in Figure 5(b), the maximum error occurs at the core and

(a) Pressure

(b) Error in Pressure

Figure 5: Contour plots of pressure and the associated RE error for an advecting vortex.

the magnitude decreases radially outward. Note that there are some small pockets of error that surround the inner core region, but these are relatively weak compared to the error computed at the centroid. These results verify that the selected error functional will direct refinement to the center of the vortex. In practice, the pressure-based error is normalized by the local pressure.

IV.B.6. Time-Accurate Behavior of Estimator

A final parameter that needs to be understood is how the applied error tolerance affects the termination refinement criterion. In other words, we seek to determine how the reported error levels change for different grid resolutions for an actual simulation. This is done by analyzing the same vortex for two time-dependent cases, one in which the vortex remains stationary, while in the other it advects across the domain. The two analogous cases are used to differentiate the impacts of error due to dissipation and the error arising from advection. Additionally, the steady case represents the best attainable error estimate, and therefore provides a manner to assess performance of the unsteady case. Both cases use five levels of refinement where $h_L = \{\frac{8}{3}, \frac{4}{3}, \frac{2}{3}, \frac{1}{3}, \frac{1}{6}\}$, allowing the coarsest and finest levels to contain about one and seventeen points across the core, respectively. Note that for these validation cases uniformly fine meshes are used, as local adaptive refinement is not applied.

For the stationary case, the vortex remains at the origin (no mean-flow) and the solution is advanced 100 coarse-grid time-steps of $\Delta t = 0.8$. The time-history of the L_∞ norm on the four finer meshes for the full simulation is presented in Figure 6. Although this is somewhat of a benign case, it highlights the ability

Figure 6: Maximum error on each level during 100 time-steps for the *stationary* vortex case.

of the higher-order finite difference solver (ARC3DC) to maintain constant error levels over the course of the simulation. Recall that at the end of each time-step, solutions on finer grids are recursively injected into their coarse equivalents. Because the vortex is well-resolved on the finest level, no observable dissipation occurs on any level, and the RE error estimator's output remains constant. Also, the error levels are increasingly spread apart for the finer meshes, demonstrating that the solution is reaching increasingly asymptotic behavior.

An earlier analysis of ARC3DC demonstrated that a resolution of about ten points across the vortex resulted in no convective error and reasonably low dispersion error⁵ when using 5th-order spatial differencing. For this test, this corresponds to the vortex on L3, which reports an error of 3×10^{-3} . This implies that an error tolerance between L3 and L4 ($3 \times 10^{-3} < e_{tol} < 3 \times 10^{-4}$) would be sufficient to maintain resolution of the vortex for this case.

For the related unsteady advecting case, it utilizes the same initial grid and solution layout, except the vortex travels along a line from $\mathbf{x}_i = \{-4, 0\}$ to $\mathbf{x}_f = \{4, 0\}$, which is a distance of about three vortex lengths. Figure 7 shows a plot of the error behavior. Note the oscillatory behavior in error that arises for the advecting case. This is because the error is measured at the grid nodes which, depending on the distance between the vortical centroid with respect to mesh nodes, causes the maximum reported error to fluctuate in a sinusoidal fashion. This is further supported by considering the L1 error that undergoes three periods during the course of the simulation, which is also equal to the number of core lengths traveled by the vortex. The same trend is observed for the L2 error, except the period is cut in half due to the Δ -to- $\frac{\Delta}{2}$ refinement strategy employed by this grid system. Despite these oscillations, the error remains periodically constant over the course of the unsteady simulation. The maximum, minimum, and average error amplitudes shown in Table 1 are comparable to the stationary case.

It should also be pointed out that the cost of applying the error estimator is quite small — about 1/150

Figure 7: Maximum error on each level during 100 time-steps for the *advecting* vortex case.

Table 1: Comparison of L_∞ error norm for a time-accurate solution of the Modified-EHS vortex when stationary or moving through the domain.

	Stationary Case	Advecting Case		
	(constant)	(min)	(max)	(avg)
L1	8.68×10^{-3}	8.63×10^{-3}	1.13×10^{-2}	1.04×10^{-2}
L2	3.10×10^{-3}	1.94×10^{-3}	3.19×10^{-3}	2.74×10^{-3}
L3	3.80×10^{-4}	3.38×10^{-4}	3.94×10^{-4}	3.75×10^{-4}
L4	1.61×10^{-5}	1.60×10^{-5}	1.66×10^{-5}	1.64×10^{-5}

the cost of a single off-body CFD iteration. The main computational cost incurred is the construction of a temporary coarse level and subsequent transfer of the solution to that level, needed for the fine-coarse error computation. Removing this overhead, the actual cost of computing the error is about one-thousandth of the time of a single CFD iteration. This is significantly more economical than an adjoint solution.

V. Integrated Feature Detection/Error Estimation Approach

The overall mesh refinement strategy which combines the feature-detection methodology with the Richardson error estimator is outlined in this section. Although it is implemented as part of *Helios*, the approach can be readily integrated into any hierarchical multi-level Cartesian-based AMR scheme.

Earlier work showed feature-detection schemes can be constructed to properly mark regions for refinement. However, since they cannot assess solution accuracy, a prescribed maximum level of resolution is applied to all features, which can severely limit overall efficiency. The local error-based approach will target refinement specifically to features with high error. Solution-based error is used to terminate refinement after a feature has been sufficiently resolved. The software package that implements this coupled error-based feature detection approach is referred to as *GAMR* (*Guided Adaptive Mesh Refinement*). Figure 8 illustrates the integration of *GAMR* into the off-body solution process of *SAMARC*. The rest of this section specifically addresses the refinement and coarsening associated with the cell identification process performed by *GAMR*.

V.A. Refinement Strategy

Suppose a sine-like feature exists on an L_N , as illustrated by Figure 9(a). Assume that L_N represents the finest level in a local region of the mesh, and that new grid levels may be only constructed atop the finest available level. For example, if a local block contained L_0 , L_1 , and L_2 grids, the only permissible way of

Figure 8: Coupling GAMR with SAMARC, the AMR-capable, off-body flow solver.

building a finer L3 grid is on top of regions of L2.

The entire wave is first identified by feature detection, as shown by Figure 9(b). Note that the feature identification process is entirely parallel, requiring only a local solution stencil for the computation of local velocity gradients. Once all features have been identified the computed RE error is used to remove refinement

Figure 9: The overall refinement process can be broken into (a) the existence of a feature on a particular grid block, (b) the initial tagging by the feature detection algorithm, (c) the removal of any tags that do not contain large error, and (d) the subsequent refinement.

in regions that are below e_{tol} . Figure 9(c) illustrates this process, as some of the previously tagged cells (green x's) are untagged (grey x's). The tagging process is now complete, and Figure 9(d) shows the selected regions receive refinement. Once the new level is created, coarse-to-fine injections initialize the solution on L_{N+1} , and the non-overlapping fine points are interpolated using second-order interpolation. Note that cells adjacent to the flagged nodes are refined, since SAMARC performs all refinement operations in a cell-wise fashion. This general process is followed for each grid level, the only exception being the coarsest level where only feature detection is applied. Since Richardson requires two overlapping solution to compute an estimate, the error on the baseline L0 grid cannot be computed.

During a regrid cycle, a maximum of one additional level can be added to an existing grid level hierarchy. Although the relative error between L_N and L_{N-1} can be used to determine if L_{N+2} grids (and finer) should be constructed, such a strategy may result in adverse consequences if the error reduction is non-linear. We prefer to have the mesh gracefully evolve with the feature, to avoid cyclical overshoot/undershoot situations.

V.B. Coarsening Strategy

For unsteady simulations, or steady flow problems with transients, coarsening operations are employed to remove refinement where no longer necessary. Recall that **SAMARC** operates a level-at-a-time, and performs regrid operations from coarse to fine. Therefore, if a feature has moved and is no longer contained in a cell situated in L_n , or if the error drops below the tolerance, *all* overlaid levels at or above L_{n+1} will automatically be removed by this process.

VI. Results

In this section, the coupled strategy is implemented within **Helios** and is allowed to drive the off-body refinement process. The test cases used in this study are: an advecting vortex, and the flow-fields generated by a NACA 0015 wing and the V-22 TRAM rotor. In each, the key objective is to verify that the feature detection, using $t_{val} = 1$, identifies the appropriate regions for refinement, and, with the help of Richardson error, the amount of applied resolution can be adjusted according to required fidelity. Particularly important for the latter practical cases, we wish to target the coherent vortical structures for refinement, rather than the error-filled turbulent wake.

VI.A. Advecting Vortex

In our previous work, the developed non-dimensional methods effectively drove the AMR process for an unsteady advecting vortex. Now, the RE error estimator is coupled with the non-dimensional Q method and is allowed to control the amount of resolution applied to the vortex. To understand the impact of different error thresholds on the time-dependent solution, three different error tolerances of 10^{-2} , 10^{-3} , and 10^{-4} are each applied to the exact same initial solution. Figure 10(a) represents the starting vortex on the six overlapping grid regions, where the grid spacing is $h_L = \{\frac{4}{3}, \frac{2}{3}, \frac{1}{3}, \frac{1}{6}, \frac{1}{12}, \frac{1}{24}\}$. The diameter of the vortex is about 3, so a single point of the coarsest mesh (L0) is contained in the core, while 72 points are found for the finest mesh (L5).

For the simulation, the vortex is placed at the origin of a domain defined by $\mathbf{x} \in [-20, 20]$ and $\mathbf{y} \in [-20, 20]$. The vortex convects rightward ($+x$) at Mach = 0.1 for a distance of 15 core widths, with periodic boundary conditions used in x , and Neumann conditions in y . A single period is $t = 300$, and by enforcing a CFL of 1.0 on the finest level, $\Delta t = 1.38 \times 10^{-2}$. For all cases, a fixed regrid period of once every thirty time-steps is employed. This frequency ensures that the mesh is adapted before the vortex is able to enter and exit an L5 cell — even if during the simulation the finer mesh levels have been removed by the adaptive error control. Therefore, during the entire 32,400 time-steps, 1,080 regrid operations are performed. Additionally, the maximum number of mesh levels is restricted from exceeding six (L5). This restriction has been imposed to provide an upper bound on the number of grid points, and since **ARC3DC** is an explicit solver, reducing the finest grid spacing would require that the time-step on the finest grid be prohibitively small.

Figures 10(b) - 10(d) contains snapshots of the final solutions when applying the three different error tolerances. Although each has travelled equal distances, as the error tolerance is reduced, additional mesh refinement is terminated earlier, which creates stronger dissipation of the vortex. For the highest tolerance of 10^{-2} , the vortex has encountered over a 65% loss of strength by the end of the simulation. When the tolerance is lowered to 10^{-3} , notable improvement is witnessed, and the final solution retains 80% of the starting maximum vorticity. This smaller tolerance has furnished an additional level of refinement, which accounts for the improvement. Finally, when using $e_{tol} = 10^{-4}$, the core is almost identical to the starting vortex, and the computed maximum vorticity has only dropped by 4%. The strong preservation of the core occurs because smallest tolerance has forced **SAMARC** to furnish three additional levels, as compared to the case where $e_{tol} = 10^{-3}$. It should be pointed out that the discontinuous contour lines do *not* represent the true solution, but are present because the visualization software encountered difficulty when computing vorticity on a multi-block grid.

To compare the time-history behavior, a measure of the actual solution error, defined as the difference between the maximum vorticity obtained from the computation and exact solution (normalized by the exact value), is plotted for the three cases in Figure 11. A value of zero represents no error, and a value of one suggests that the vortex has been entirely destroyed. Beginning with the smallest tolerance, a value of 10^{-4} results in low error levels that grow logarithmically for the entire simulation. Next, for $e_{tol} = 10^{-3}$, larger, but reasonable, error growth is obtained. The curve appears to be nearly linear, suggesting that the

(a) Initial Solution. Note: L1 and L0 surround the L2 grid system, but are not represented above.

(b) Final Solution, $e_{tol} = 10^{-2}$

(c) Final Solution, $e_{tol} = 10^{-3}$

(d) Final Solution, $e_{tol} = 10^{-4}$

Figure 10: The initial vortex solution is created on six levels of resolution, and the RE error estimator is allowed to add or remove resolution if and when required at each adapt step. The finest resolution is capped at L5 and levels are removed if the computed error is less than e_{tol} .

error will steadily rise if the vortex is allowed to continue. Finally, when the tolerance is raised to 10^{-2} , significantly higher error occurs, largely due to the initial jump. The spike, which happens during the first core width, demonstrates how important *maintaining* sufficient resolution is to vortex preservation. It is, however, interesting to note that after the initial loss of 30%, the error rate lessens and gradually advances toward an upper limit of about 70%. This trend illustrates that as the vortex continually dissipates, which causes an enlargement of the core region, it will reach a point where the underlying grid can provide adequate resolution. Once this occurs, the core growth, and therefore the associated error, will increase at much slower rate.

Table 2 compares the number of grid points required by each applied error tolerance, and the corresponding final error reported. Reducing the error tolerance from 10^{-2} to 10^{-3} results in an increase of nodes by a factor of 1.7 and a decrease in error by 3.5 times. Further reducing the error tolerance to 10^{-4} provides diminishing returns, as the relative error decreases by a factor of five, but requires 26 times more grid points. While the level of acceptable solution fidelity is dependent upon the particular application, as demonstrated when using a tolerance of 10^{-3} , an optimal e_{tol} is likely to exist which substantially reduces solution error

Figure 11: Error in vorticity with respect to distance travelled by the vortex on an adaptive grid system whose resolution is controlled by different error tolerances.

Table 2: The number of required grid points and the final reported error, when using various error tolerances to control grid resolution.

e_{tol}	Finest Level	# of Nodes	Final Error
10^{-2}	L1	1.2×10^5	67%
10^{-3}	L2	1.7×10^5	20%
10^{-4}	L5	4.5×10^6	4%

at a relatively marginal cost.

VI.B. NACA 0015 Wing

This test case is based on the original experiment of McAlister and Takahasi,³⁰ where a series of vortex core measurements were taken at several downstream locations for different angles of attack, freestream velocities, and wing tip shapes. The current test case considers the steady flow around a full-span NACA 0015 square wing at a 12° angle of attack. It experiences uniform inflow of 0.1235 Mach with a Reynolds number of 1.5×10^6 . Dirichlet conditions are imposed along the far-field boundary. The **Helios** dual-mesh infrastructure, outlined by Section II, is employed and the hybrid unstructured/structured-Cartesian grid system at the start of the simulation (before any solution-based refinement) is illustrated in Figure 12. The unstructured grid surrounds the wing section and is embedded within the Cartesian mesh. The off-body domain pictured here is comprised of four mesh levels where $h_{L0} = 0.2$, $h_{L1} = 0.1$, $h_{L2} = 0.05$, and $h_{L3} = 0.025$, where the grid distance has been non-dimensionalized by the chord length. In general, regridding was performed every 100 iterations, and convergence was obtained after about 250 adaption steps.

In previous work, this case was run with feature-only adaption and the improved solution underscored the effectiveness of this type of refinement for vortex preservation. The main conclusion is re-presented in Figure 13, which compares two converged solutions — one with, and one without solution adaption. The use of the fine mesh system that surrounds the trailing vortex allows it to remain coherent, even after nearly twenty chord lengths downstream. Furthermore, these tests showed that non-dimensional feature detection algorithms were able to call for appropriate regions of refinement when using threshold values of one. Although these results were adequate, core-velocity comparisons to experimental data revealed that the computed vortices undershot experimental values by about 40% on average, and preliminary results suggested that the unstructured mesh was the root cause.

In an effort to understand the role that grid-dependence plays on the dissipation of vortical features, a new near-body grid with enhanced resolution near the tip is applied. Since the dual-mesh grid system is not

Figure 12: The NACA 0015 near-body unstructured grid with the unadapted off-body Cartesian grid system.

Figure 13: Grid systems and isosurfaces of vorticity magnitude (drawn at the same value) when solution-based refinement is, and is not, applied for the NACA 0015. On the right, the non-dimensional Q method is allowed to control the refinement process using a threshold value of unity.

only defined by two unique grid systems, i.e. unstructured near-body and Cartesian off-body, but also the distance at which the near-body extends into the far-field, three different studies are conducted. Starting at the wing surface, the near-body grid is replaced with another that offers significantly more refinement near the wing tips. Next, the off-body resolution is increased to five levels, where a refined L4 mesh surrounds the wing and any identified features. For the last grid study, the overlap region is adjusted so that the near-body grid is trimmed inward, allowing the trailing vortex a shorter distance to travel before it reaches the off-body grid. After the grid studies are completed, error-adaptive control is allowed to adjust resolution for a setup that demonstrates reasonable performance.

VI.B.1. Near-Body Grid Resolution

Since it was discovered that the “original” mesh (Figure 14(a)) caused dissipation of vorticity due to coarseness near the tip, a “tip-refined” mesh (Figure 14(b)), which has a concentration of points at the wing

tips, is tested and compared against the original mesh. Note that the tip-refined mesh is not generated by

Figure 14: Identical y -cut planes of the two different unstructured grids tested. The grid slice rests at the tip of the wing to juxtapose the regions with and without additional tip refinement.

adding resolution to the original mesh as indicated by its reduced nodal count, but was created anew. To compare the two solutions, z -velocities (normalized by freestream) are measured through the vortex core, at one and six chord lengths downstream, to coincide with experimental data. The 2D line plots are created by a spanwise sampling through the vortex core, as indicated by Figure 15.

Figure 15: Vortical strength is measured at two downstream locations by extracting a spanwise line, which passes through the core, of z -velocity data (normalized by the freestream).

Surprisingly, Figure 16 suggests that the overall dissipation reported in the off-body is roughly equivalent when using either near-body mesh. At one chord back, the computed solutions show a 40% loss in maximum z -velocity, as compared to experiment. Although at six chords downstream, the discrepancy between the experimental and computed solutions is reduced, the two meshes produce nearly identical results. These results indicate that a finer near-body mesh cannot singlehandedly improve vortex prediction in the far-field. For that reason, we next investigate the effects of increasing off-body resolution.

Figure 16: A comparison of normalized z-velocity recorded at one and six chords downstream for the original and tip-refined meshes when four levels of refinement are used in the off-body domain.

VI.B.2. Off-Body Resolution

The original off-body grid system, shown in Figure 12, is refined so that the unstructured mesh is now surrounded by an L4 mesh. Since each refined level halves the grid spacing of its coarse equivalent, $h_{L4} = 0.0125c$. To be clear, not only is the refinement near the unstructured grid increased, but any feature identified in the off-body will also receive this finer resolution. Figure 17 separately compares the effect on vortex preservation for the original and tip-refined meshes at six chord lengths downstream. While a

Figure 17: Normalized z-velocity at six chords downstream for the original and tip-refined meshes, when using an additional level of off-body refinement. For the original mesh, the maximum value is increased by 17%, and is increased by 31% for the tip-refined mesh.

notable increase of velocity is witnessed for both cases, the improvement shown by the tip-refined mesh (31%) is about twice as large as found for the original mesh (17%). Therefore, only after adding refinement in the off-body domain, are the advantages of proper near-body resolution be realized. The coupled effects between the near- and off-body grids is not surprising, as the key to vortex preservation relies upon the adequate resolution of both. However, since the off-body solver is more efficient, due to the structured grid, and accurate, due to the fifth- versus second-order accurate spatial discretizations, it is desirable to have the vortex enter the off-body grid system as soon as possible. Therefore, as a final grid study, the near-body

grid is trimmed inward.

VI.B.3. Near-Body Cut Distance

The cut distance refers to the minimum length at which the full unstructured mesh is trimmed inward toward the wing. Figure 18 illustrates cut distances of $0.25c$ and $0.1c$ for the original mesh when using five levels of off-body refinement. Since the off-body mesh is automatically generated on startup, it automatically fills the

Figure 18: A side-by-side comparison of the two near-body grid cut distances used, here shown on the original NACA unstructured mesh with five level of off-body refinement.

displaced unstructured region with a structured grid that sits closer to the wing. Figure 19 illustrates the effects of trimming the original and tip-refined meshes by considering vortical strengths at six chord lengths downstream. For the original mesh, a smaller near-body grid results in a 14% improvement. However, when

Figure 19: The effects of trimming the near-body grid, from a quarter- to a tenth-chord, on the trailing vortex at six chord lengths back. Five levels of off-body resolution are employed in both cases.

the tip-refined mesh is trimmed, no notable difference is reported. Interestingly, the tip-refined maximum strengths now match that which is obtained for the trimmed original mesh.

These three grid-based studies confirm that, to effectively reduce the dissipation of vortical structures in the far-field, the improved near-body resolution can help, but only if resolution is comparable. The off-body grid can be used to maintain strength over long distances, but only if the feature can pass through the near-body region without experiencing significant dissipation. Furthermore, as demonstrated by the original mesh, if the resolution of the unstructured grid is too coarse, trimming the mesh can help alleviate the effects of numerical dissipation in the far-field.

VI.B.4. Error-Adaptive Control

In this section, error-based adaption is applied in conjunction with feature detection to control the amount of applied off-body resolution. Specifically, error tolerances of 10^{-2} and 10^{-3} are used to again drive refinement for the NACA 0015 case. Here, the original mesh is used with a cut distance of $0.25c$, and the off-body grid contains a maximum of five levels of refinement. Figure 20 compares the different solutions that are created when using the different e_{tol} thresholds, where higher error tolerances can effectively limit the amount of resolution applied. Setting $e_{tol} = 10^{-2}$ only allows the vortex to travel about three chord lengths downstream,

Figure 20: Adaptive grids and isosurfaces of vorticity magnitude for the NACA 0015 case when two different error tolerances are used to adjust grid resolution in regions where the non-dimensional Q method has allowed for refinement.

while a threshold of 10^{-3} allows them to reach the end of the computational domain. A closer comparison of the solutions obtained is given in Table 3. The prior case, which was run without error-based control, is used as a benchmark, and relative error for each case is computed at 1, 6, and 12 chord lengths downstream. An e_{tol} of 10^{-3} results in perfect agreement, while the larger tolerance is seen to be substantially more

Table 3: The effect upon solution accuracy that different error tolerances have for the NACA 0015 flow-field. Error comparisons are made to the solution that was obtained using a feature-only mode, where any feature found received maximum refinement.

	# Nodes	Relative Error		
		at 1 chord	at 6 chords	at 12 chords
$e_{tol} = 10^{-2}$	12.1×10^6	3%	57%	68%
$e_{tol} = 10^{-3}$	19.8×10^6	0%	0%	0%

dissipative. Although at one chord back, setting $e_{tol} = 10^{-2}$ only results in a 3% loss in strength, a severe increase occurs over the next five chord lengths. The effects of numerical dissipation are less between 6 and 12 chords. Note that the rapid increase, followed by moderate growth, mimics the behavior as previously witnessed for the advecting vortex when under-resolved.

Examining the differences between supplied grid refinement, Figure 21 provides a close-up view of the grids generated by both error tolerances just behind the wing tip. With the larger error tolerance, a telescoping-grid effect is witnessed, whereby the finer mesh levels are removed just past the trailing edge. And although this tolerance still furnishes the maximum level of refinement, the region to which it is applied is rather small. In contrast, when the tolerance is lowered to 10^{-3} , the fine mesh region grows so that the trailing vortex core is fully enveloped.

Figure 21: A close-up the off-body adaptive grid systems without the computed solution and with the NACA 0015 wing. The grid has been adapted by using RE error to control the degree of refinement that the trailing vortex cores receive.

Although these plots clearly demonstrate that an RE-based estimator is capable of controlling the amount of resolution that a vortical feature receives, the significance of the feature detection algorithm should not be overlooked. By only allowing the RE estimator to refine where features are found, and *not* simply to regions with high error, the current approach avoids refining the turbulent wake. Figure 22 shows that by

Figure 22: A cutting plane of the adaptive mesh containing a contour plot of error, located about half a chord downstream from the trailing edge of the NACA 0015 wing, confirms that high regions of error are found in the wake and the vortex. However, by only permitting refinement of coherent structures, no attempt is made to resolve the wake.

specifically targeting the vortex cores, refinement is not wasted on the wake, but is directed towards the vortex cores.

VI.C. Tilt Rotor Acoustic Model (TRAM)

The Tilt Rotor Acoustic Model (TRAM) is a quarter-scale model of the Bell/Boeing V-22 Osprey tiltrotor aircraft right-hand 3-bladed rotor. The isolated TRAM rotor was tested in the Duits-Nederlandse Wind-tunnel Large Low-speed Facility (DNW-LLF) in the spring of 1998. Aerodynamic pressures, thrust and

power, were measured along with structural loads and aeroacoustics data. Wake geometry, in particular the locations of tip vortices, was not part of the data collected. Further details on the TRAM experiment and extensive CFD validations with OVERFLOW can be found in the work of Potsdam and Strawn.²

For the computational tests, $M_{tip} = 0.625$, $Re_{tip} = 2.1 \times 10^6$, and the collective is set to 14° . A non-inertial reference frame is used, such that the rotor stays fixed within a rotational freestream set with moving grid source terms. Although the freestream Mach number is low, the speed of the flow relative to the blade is high due to the rotational terms, so low-Mach preconditioning is not applied. The TRAM geometry contains multiple components, the 3 blades and a center-body. The near-body volume mesh contains 8.3×10^6 nodes, and extends approximately two chord lengths from the blade surface. It contains relatively fine resolution at the tips and under the blade to improve vortex generation, and to reduce vortex dissipation, as retreating blades pass over the trailing vortices. The maximum number of off-body levels is seven, which corresponds to a length that is 10% of the rotor tip chord. The unstructured grid, as it is subset within the unadapted Cartesian mesh, is given in Figure 23.

Figure 23: The TRAM near-body unstructured grid with the unadapted off-body Cartesian grid system.

VI.C.1. Automated Feature-Based Refinement

In this section, we apply the non-dimensional Q method, initially without adaptive error control, to steer the AMR process. A feature-only approach is first considered as non-dimensional feature detection algorithms have never driven the AMR process for this case. A similar study was performed by Wissink³¹ *et al.*, but off-body refinement was controlled by vorticity magnitude, which required that the threshold be specifically tuned for this particular problem. Here, however, the non-dimensional Q method uses a threshold of one and is allowed to mark regions for refinement. To provide a baseline comparison, Figure 24 plots a case without refinement alongside another when solution-based refinement is employed. It is obvious that without additional refinement, the trailing vortices experience strong dissipation and only remain coherent for only about a sixth of a revolution. In contrast, by providing adequate off-body resolution, the case with adaption allows the vortical structures to remain coherent, preserving their strength for several blade revolutions. Table 4 offers a quantitative comparison of the rotor thrust and torque coefficients (C_T , C_Q), and the figures of merit ($FM = C_T^{3/2}/C_Q\sqrt{2}$) obtained from experiment and by simulation. Although C_T and C_Q predictions

Table 4: Experimental and computed rotor thrust and torque coefficients (C_T , C_Q), the figures of merit (FM), and the associated error with and without solution-based adaption in the off-body mesh.

	C_T	C_Q	FM	FM_{error}
Experiment	0.0149	0.00165	0.779	–
No Adaption	0.0147	0.00172	0.736	5.5%
With Adaption	0.0151	0.00171	0.771	1.0%

Figure 24: TRAM solutions generated with and without solution-based adaption when the non-dimensional Q method is used to guide the AMR process. A y -cut plane depicts the off-body grid, and the isosurfaces shown represent equivalent values of Q , which are colored by vorticity magnitude.

are only marginally better when using adaption, the figure of merit significantly improves. This metric, which is a measure of the rotor's overall efficiency, is seen to benefit from the highly resolved wake. Further, it should be noted that the tabulated data found here when using a locally-normalized feature detection algorithm is identical to the results originally obtained by Wissink *et al.*

VI.C.2. Error-Adaptive Control

We next test the ability of the RE error estimator to control the AMR process for the TRAM case. The maximum number of off-body grid levels is still capped at seven, but the adaptive error-based control will turn off adaptivity on features where the computed error is below a threshold e_{tol} . A coarser near-body grid (4.9M nodes) is used for these tests.

Two different error tolerances are applied, $e_{tol} = 10^{-4}$ and $e_{tol} = 10^{-5}$. The resulting flowfields and adaptive Cartesian grid systems in Figure 25. While the developed normalized pressure functional is not

Figure 25: Isosurfaces of Q , colored by vorticity magnitude, overlaid on the corresponding adaptive Cartesian grid system generated for the TRAM case when two different error tolerances are used.

entirely problem independent, the results clearly indicate that the RE error estimator can effectively control

resolution. The amount of applied resolution can be effectively controlled by the imposed error tolerance. With $e_{tol} = 10^{-4}$ the estimator prevents the finest allowable grid level from being used, as six levels of refinement proves to be adequate in the near-field. By following pathlines generated by the inner and outer helical vortex systems, one notices that the amount of refinement is continually reduced, which leads to stronger dissipative effects and therefore a shrinking of the cores. Note that the inner vortex system, generated by the inboard rotor tips, is considerably stronger than the outer tip vortices. This results in higher discretization error, which creates additional refinement for this region. While the outer tip cores are still detectable, the inadequate resolution makes them decay at a relatively faster rate. However, when $e_{tol} = 10^{-5}$, the complete vortical system is preserved for significantly longer. The case-maximum seven levels of refinement is applied near the rotor. Note that the adaptive error-based control is still active, as witnessed by the fine-grid region only extending for about 1.5 revolutions.

While additional studies are necessary to understand the effects of different error thresholds upon rotor performance and wake prediction, here it is verified that the developed AMR procedure can effectively be used to guide refinement so that a desired level of fidelity is achieved.

VII. Conclusion

An AMR strategy is developed for unsteady vortex-dominated flows which couples feature detection with Richardson-extrapolation (RE) error detection to control grid resolution. The paper focuses mainly on the development of the RE error estimator, the feature detection algorithms used have been presented in earlier work.^{4,5}

The error estimator is validated with a grid convergence study using vortices of various resolution advecting a range of distances. Spatial and temporal error was determined by comparing to an exact known solution. The RE scheme was found to converge asymptotically for such flows in the limit as the mesh spacing approaches zero, showing ideal logarithmic convergence rates that matched the *LTE* and removed the effects of temporal error.

The RE estimator was used in conjunction with the non-dimensional Q feature detection method to control the off-body AMR process. The capability of the approach was tested for three cases, an advecting vortex, tip vortices from a NACA 0015 wing, and the wake of a quarter-scale V22 TRAM rotor. For the vortex case an error tolerance of 10^{-4} was found to preserve the vortex core for more than 15 core widths. For the NACA0015 wing, a tolerance of 10^{-3} resulted in complete refinement of the tip vortices 20 chords behind the wing toward the back of the domain. The use of the feature detection in the approach directed refinement to the vortex cores and not to broad regions in the wake that contained high error. The scaled V-22 TRAM rotor, which was studied using off-body refinement controlled by feature detection alone as well as in conjunction with the error estimator, was found to properly resolve the feature-rich wake. The off-body refinement improved rotor performance metrics and reduced the figure of merit error from 5% to 1%, compared to a case with no adaption. The error control demonstrated effectiveness in controlling the off-body refinement.

Future work will apply the developed approach for more complex rotorcraft applications, in particular those involving multiple components such as fuselage and tail rotor interaction.

Acknowledgments

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. 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. 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

¹Pulliam, T. H. and Jespersen, D. C., "Large Scale Aerodynamic Calculation on Pleiades," 21st International Conference on Parallel Computational Fluids Dynamics, May 2009, Moffett Field, CA.

- ²Potsdam, M. A. and Strawn, R. C., "CFD Simulations of Tiltrotor Configurations in Hover," 58th Annual American Helicopter Society Forum, Jun. 2002.
- ³Caradonna, F. X., "Developments and Challenges in Rotorcraft Aerodynamics," 38th AIAA Aerosciences Conference, Jan. 2000, AIAA Paper 2000-0109.
- ⁴Kamkar, 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.
- ⁵Kamkar, S. J., Jameson, A. J., Wissink, A. M., and Sankaran, V., "Feature-Driven Cartesian Adaptive Mesh Refinement in the Helios Code," 48th AIAA Aerosciences Conference, Jan. 2010, AIAA Paper 2010-0171.
- ⁶Hunt, 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.
- ⁷Jeong, J. and Hussain, F., "On the Identification of a Vortex," Journal of Fluid Mechanics, Vol. 285, 1995, pp. 69–94.
- ⁸Chong, 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.
- ⁹Horiuti, K. and Takagi, Y., "Identification Method for Vortex Sheet Structures in Turbulent Flows," Physics of Fluids, Vol. 17, No. 12, Dec. 2005.
- ¹⁰Venditti, D. A. and Darmofal, D. L., "A Multilevel Error Estimation and Grid Adaptive Strategy for Improving the Accuracy of Integral Outputs," 37th AIAA Aerosciences Conference, Jan. 2008, AIAA Paper 1999-3292.
- ¹¹Nemec, 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.
- ¹²Duraisamy, K., Alonso, J. J., and Chandrasekhar, P., "Error Estimation for High Speed Flows using Continuous and Discrete Adjoint," 48th AIAA Aerospace Sciences Meeting, Jan. 2010, AIAA Paper 2010-128.
- ¹³Krakos, J. A. and Darmofal, D. L., "Effect of Small-Scale Unsteadiness on Adjoint-Based Output Sensitivity," 19th AIAA Computational Fluid Dynamics Conference, 2009, AIAA Paper 2009-4272.
- ¹⁴Sankaran, V., Sitaraman, J., Wissink, A. M., Datta, A., Jayaraman, B., Potsdam, M., Mavriplis, D., Yang, Z., O'Brien, D., Saberi, H., Cheng, R., Hariharan, N., and Strawn, R., "Application of the Helios Computational Platform to Rotorcraft Flowfields," 48th Aerospace Sciences Meeting and Exhibit, Jan. 2010, AIAA Paper 2010-1230.
- ¹⁵Wissink, 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, Jan. 2008, AIAA Paper 2008-927.
- ¹⁶Sitaraman, 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.
- ¹⁷Jespersen, D., Pulliam, T., and Buning, P., "Recent Enhancements to OVERFLOW," 35th AIAA Aerosciences Conference, Jan. 1997, AIAA Paper 1997-0644.
- ¹⁸Buning, P., Gomez, R., and Scallion, W., "CFD Approaches for Simulation of Wing-Body Stage Separation," 22nd AIAA Applied Aerodynamics Conference, Aug 2004, AIAA Paper 2004-4838.
- ¹⁹Wissink, A. M., Kamkar, S. J., Sitaraman, J., and Sankaran, V., "Cartesian Adaptive Mesh Refinement for Rotorcraft Wake Resolution," 28th Applied Aerodynamics Conference, Jun. 2010.
- ²⁰Hornung, 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.
- ²¹Hornung, 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.
- ²²Wissink, A. M., Hornung, R. D., Kohn, S. R., Smith, 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.
- ²³Berger, M. and Oliger, J., "Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations," Journal of Computational Physics, Vol. 53, 1984.
- ²⁴Berger, M. J. and Colella, P., "Local Adaptive Mesh Refinement for Shock Hydrodynamics," Journal of Computational Physics, Vol. 82, May 1989, pp. 64–84.
- ²⁵Wissink, A. M., Hysom, D. A., and Hornung, R. D., "Enhancing Scalability of Parallel Structured AMR Calculations," Proceedings of the 17th ACM International Conference on Supercomputing (ICS03), Jun. 2003, pp. 336–347.
- ²⁶Gunney, B. T. N., Wissink, A. M., and Hysom, D. A., "Parallel Clustering Algorithms for Structured AMR," J. Parallel Distrib. Comput., Vol. 66, No. 11, 2006, pp. 1419–1430.
- ²⁷Kamkar, S. J., Mesh Adaption Strategies for Vortex-Dominated Flows, Ph.D. thesis, Stanford University, Expected Date of Publication: Feb. 2011.
- ²⁸Erlebacher, G., Hussaini, M. Y., and Shu, C.-W., "Interaction of a Shock with a Longitudinal Vortex," Journal of Fluid Mechanics, Vol. 337, 1997, pp. 129–153.
- ²⁹Roy, C. J., "Review of Discretization Error Estimators in Scientific Computing," 48th AIAA Aerosciences Conference, Jan. 2010, AIAA Paper 2010-126.
- ³⁰McAlister, K. W. and Takahasi, R. K., "NACA 0015 Wing Pressure and Trailing Vortex Measurements," NASA Technical Paper, Vol. 3151, Nov. 1991.
- ³¹Wissink, A. M., Potsdam, M., Sankaran, V., Sitaraman, J., Yang, Z., and Mavriplis, D., "A Coupled Unstructured-Adaptive Cartesian CFD Approach for Hover Prediction," American Helicopter Society 66th Annual Forum, May 2010.