

ACCURATE AND FLEXIBLE FORMULATION OF A DUAL-SOLVER HYBRID CFD FRAMEWORK

Alex M. Moushegian, amoushegian3@gatech.edu, Georgia Institute of Technology (USA)
Marilyn J. Smith, marilyn.smith@ae.gatech.edu, Georgia Institute of Technology (USA)
Glen R. Whitehouse, glen@continuum-dynamics.com, Continuum Dynamics, Inc. (USA)
Dan A. Wachspress, dan@continuum-dynamics.com, Continuum Dynamics, Inc. (USA)

Various hybrid computational fluid dynamics (CFD) methods have recently been developed which couple Navier-Stokes solvers to vortex-based solvers for cost-effective rotorcraft aeromechanical analysis. Previous hybrid methods have had implementation issues which result in inaccurate and/or inflexible frameworks. Methodological approaches to CFD/free-wake coupling that address these issues and are unique among other hybrid CFD methods are introduced, including rotor tracking, blade deformation treatment, unsteady free-wake boundary pressure, boundary characteristic treatment, boundary value interpolation, and single gridded blade (SGB) simulation. These improvements are implemented into the OVERFLOW-CHARM hybrid CFD framework and their impact is demonstrated using three example scenarios: vortex advection, wing-integrated propulsion, and a rotor in hover. The new methods quantifiably improve the quality of the hybrid CFD solutions in these cases and enhance the generality of the framework to problems related to future vertical lift.

1. Introduction

Modern advances in computing hardware and software have allowed computational methods for the prediction of rotorcraft aerodynamics to become powerful and invaluable tools for rotorcraft design and analysis. Current computational methods for rotorcraft aeromechanical analysis can generally be grouped into high-fidelity computational fluid dynamics (CFD) methods, mid-fidelity vortex-based wake codes, or lower-fidelity momentum-based aerodynamic codes. While CFD methods can provide accurate aerodynamic analysis derived from first principles, they are typically too expensive to be used early in the design process or for repeated analysis. In these situations, mid-fidelity aerodynamic codes are used to generate acceptable results given time or resource constraints, but make various assumptions which limit the models' accuracy when applied to complex or novel configurations. Momentum-based models are typically used for very early design calculations or as a component of more advanced models.

2. Current Hybrid CFD Methods

To bridge the gap between high- and mid-fidelity codes, hybrid CFD methods have been developed in the last two decades which couple CFD solvers to secondary aerodynamic codes for rotorcraft analysis. Depending on the secondary code, these hybrid solvers can provide various amounts of computational cost savings with different levels of fidelity. These benefits also tend to vary with the specific application, though previously-developed methods have had limitations inherent to their formulation.

For example, the hybrid solver presented by Anusonti [1] couples RANS CFD to a Particle Vortex Transport Method (PVTM). The boundary condition for the RANS domain is set to the flow conditions induced by the PVTM solution. The RANS flow field is converted to vortex particles in the PVTM simulation. Depending on the resolution of the PVTM simulation, this method can conserve vorticity at reduced cost com-

pared to a full RANS simulation, and can represent arbitrary vorticity fields which may be difficult to represent with more structured vorticity elements, such as vortex filaments or a vortex lattice. However, representing large, dense, and well-ordered vorticity fields like wing and rotor wakes with vortex particles is less computationally efficient than using a structured vortex method. The poor scaling of PVTM solvers with the number of particles exacerbates this issue. The implementation of the boundary condition also assumes free-stream density everywhere, which may introduce errors near the rotor tip where compressibility effects should not be ignored. The pressure is also purely extrapolated from the CFD domain, with no direct influence from the PVTM solution.

Another example is the GT-Hybrid methodology, which has coupled both a free-wake lattice and PVTM simulation with an in-house CFD solver to represent rotor wakes in various configurations [2, 3]. The advantages of each potential solver have not been quantifiably established via any publications to date. It also employs an inaccurate approach to the boundary condition treatment for hybrid CFD/free-wake simulations. Free-wake elements within the CFD domain are not included when computing the induced conditions on the domain boundary in an attempt to prevent double-counting of vorticity predicted by both the CFD and free-wake solutions. It also employs a characteristic boundary condition based on Riemann invariants to mitigate reflection of outgoing waves. However, these two decisions mean that inflow boundary conditions are not influenced by the rotor wake within the CFD domain, since the boundary conditions are set to the conditions induced solely by the free-wake solution, which is neglected within the CFD domain. Also, outflow boundary conditions do not include the influence of the rotor wake outside the CFD domain, since the boundary conditions are extrapolated from the interior CFD solution. This will result in non-physical or inaccurate conditions at the CFD domain boundary [4].

Finally, the hybrid coupling between the CFD solver OVERTURNS and free-tip-vortex wake solver PWAM presented by Thomas [5] couples the two solvers using a field-velocity approach, which avoids reflection issues at the boundary, but introduces double counting of velocity induced by vorticity within the CFD domain. This approach is also significantly more expensive than boundary-condition-based hybrid methods, as induced velocities have to be computed at every point in the CFD grid instead of just at the boundary.

The problems with current hybrid CFD methods can be summarized as follows. First, the solvers employ in-house codes which tend to be poorly supported and unwieldy to new users. Second, the codes are gen-

erally tailored to specific applications to demonstrate the merits of the basic premise, and are thus inflexible when applied to complex configurations and scenarios. Finally, they make certain implementation decisions which may limit their accuracy or computational efficiency. These limitations hinder the application of previous hybrid CFD solvers to current advanced vertical takeoff and landing (VTOL) designs.

3. Current Methodology

3.1. Basic Formulation

The hybrid CFD code OVERFLOW-CHARM has the basic advantage that it couples component codes that are widely used in the rotorcraft community and have been validated for many complex configurations and scenarios. The CFD component is OVERFLOW, developed and supported by the National Aeronautics and Space Administration (NASA) and the US Army. OVERFLOW is a structured overset unsteady Reynolds-Averaged Navier-Stokes (URANS) solver capable of resolving detailed flow physics around complex geometries [6]. The structured grid format permits the use of high order numerical schemes for reduced numerical dissipation of vorticity in the flow field. The free-wake component is CHARM, developed and supported by Continuum Dynamics, Inc. (CDI). CHARM is a Lagrangian full-span free-vortex wake solver which can efficiently compute rotorcraft wake aerodynamics, including wake-body interactions [7, 8]. It uses curved vortex elements that allow complex rotor wakes to be resolved using fewer elements than would be required when using linear elements.

The coupling framework of OVERFLOW-CHARM is described in Figure 1. The OVERFLOW and CHARM solutions communicate to each other through the OVERFLOW-computed blade loads and the CHARM-computed CFD domain boundary flow conditions. This method allows the CFD domain to be significantly reduced compared to a conventional CFD simulation while still accounting for the effects of the full rotor wake. The original formulation was similar to that of GT-Hybrid, where wake-induced conditions are added to the freestream conditions; pressure and density are computed using isentropic relations; and A characteristic condition was then applied to the boundary. The basic framework, without the recent improvements discussed in this work, was demonstrated in previous studies to provide accurate rotor performance metrics and dynamic loads within 5% of conventional CFD methods (OVERFLOW) at less than half the cost for rotors in hover and forward flight [9, 10, 11].

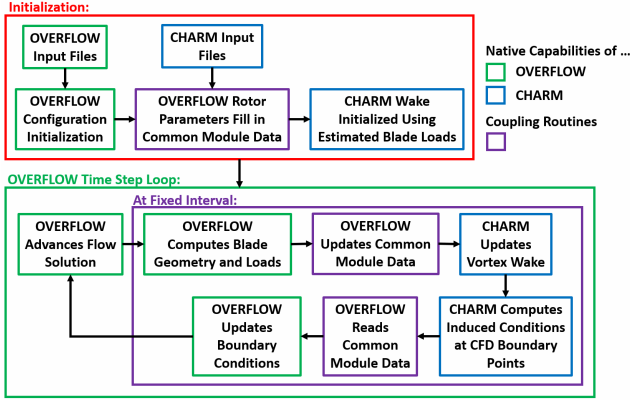


Figure 1: Coupling framework of OVERFLOW-CHARM

4. Recent Improvements to OVERFLOW-CHARM

4.1. Advanced Rotor Configurations

The OVERFLOW-CHARM framework was originally unable to model multiple rotors or rotors with hub transformations relative to the inertial frame. With the push toward advanced future vertical lift designs, it is important for a rotorcraft analysis tool to be able to handle complex rotor-rotor and rotor-fuselage configurations. Both OVERFLOW and CHARM have native capability for modeling multiple rotors. OVERFLOW includes the number of rotors as a user input, as well as each rotor’s reference conditions, number of blades, and CFD grids associated with each blade. The blade geometry is derived from the surface grid topology, and blade deformations are driven either by coupling to a structural solver or with prescribed blade motion. CHARM requires all rotor and blade properties as inputs. The coupling routines between OVERFLOW and CHARM were re-written such that all the required information was passed from OVERFLOW to CHARM, regardless of the complexity of the configuration. This enables simple configurations like a rotor-fuselage case with a rotor shaft angle, as well as more complex configurations like coaxial or tandem rotor cases up to the potential for full-vehicle analysis with any number of rotors to be modeled, while reducing user input errors.

4.2. Blade Deformation Treatment

Previously, the OVERFLOW-CHARM framework required a specific blade deformation file as an input, even when running with rigid blades. This also limited coupling of OVERFLOW-CHARM to particular

structural solvers which perform CFD-CSD coupling through file I/O. To address this, the OVERFLOW-CHARM coupling routines were refactored to be agnostic to CFD-CSD coupling information and to compute blade geometry directly from the deformed CFD grid. This removes the requirement of a null blade motion file when rigid motion is desired. It also permits CFD-CSD coupling between OVERFLOW-CHARM and RCAS through the HPCMP CREATE-AV™ Helios framework, which deforms the CFD grid through the Melodi mesh manager rather than through a blade motion file.

4.3. Boundary Condition Improvements

4.3.1. Unsteady Free-Wake Pressure

To improve the accuracy of OVERFLOW-CHARM when the CFD domain is reduced to a region very close to the rotor blades, it was important to develop a robust and physically consistent CFD boundary condition so that the CHARM solution accurately influences the OVERFLOW solution. Previously, the far field boundary condition in OVERFLOW-CHARM only applied the instantaneous induced velocities from the CHARM solution to update the boundary velocities and extrapolated the density and pressure using isentropic relations. Pressure and density are now computed using the unsteady potential function $\frac{\partial\phi}{\partial t}$ predicted by CHARM as follows:

$$P = \frac{1}{2}\rho_{\infty} \left(U_{\infty}^2 - U_{induced}^2 - 2\frac{\partial\phi}{\partial t} \right) + P_{\infty} \quad (1)$$

$$\rho = \rho_{\infty} + \frac{1}{a_{\infty}^2} (P - P_{\infty}). \quad (2)$$

More details on the unsteady potential function in CHARM can be found in Wachspress et. al [12]. This improves the accuracy of the boundary condition when the boundary is immersed in a highly unsteady portion of the flow field and is critically important for computing BVI loads on a fuselage. This situation is common for the most economical set-up of an OVERFLOW-CHARM case, so tracking the unsteady pressure means that accuracy is maintained while providing significant computational cost savings.

4.3.2. Free-Wake Characteristic Boundary Condition

In previous implementations, a standard preconditioned characteristic boundary condition (CBC) based on Riemann invariants (henceforth referred to as a standard CBC) was applied after setting the bound-

ary values to the CHARM-induced conditions. This mimics the conventional far field boundary treatment in OVERFLOW. The algorithm for setting boundary values $\vec{Q}^{(r)}$ based on internal point values $\vec{Q}^{(i)}$ and external point values $\vec{Q}^{(o)}$ is as follows:

$$\vec{Q}^{(r)} = \vec{\bar{Q}} - \frac{1}{2} \mathbf{M} \mathbf{X} \text{sign}(\mathbf{\Lambda}) \mathbf{X}^{-1} \mathbf{M}^{-1} \Delta \vec{Q}, \quad (3)$$

where

$$\vec{\bar{Q}} = \frac{\vec{Q}^{(o)} + \vec{Q}^{(i)}}{2}, \quad (4)$$

$$\Delta \vec{Q} = \frac{\vec{Q}^{(o)} - \vec{Q}^{(i)}}{2}. \quad (5)$$

\mathbf{M} , \mathbf{X} and the sign function on $\mathbf{\Lambda}$ together treat the preconditioned eigenvalues appropriately depending on whether the boundary sees subsonic/supersonic inflow/outflow and compute the effect on the primitive variables \vec{Q} . The preconditioned eigenvalues of the system are:

$$\lambda_1 = \vec{u} \cdot \hat{n}, \quad (6)$$

$$\lambda_2 = \frac{1}{2} (\vec{u} \cdot \hat{n}) (1 + \beta) + s_k, \quad (7)$$

$$\lambda_3 = \frac{1}{2} (\vec{u} \cdot \hat{n}) (1 + \beta) - s_k, \quad (8)$$

where \hat{n} is the Cartesian vector normal to the boundary face, s_k is analogous to the speed of sound, and β determines the strength of the preconditioning.

The motivation for this boundary condition is to prevent reflections of outgoing characteristics back into the CFD domain. These characteristics can be split into advective (λ_1) and pressure wave (λ_2, λ_3) characteristics. Taking the sign of $\mathbf{\Lambda}$ serves as a test for which characteristics are leaving the CFD domain and to treat them appropriately to prevent their reflection.

This formulation of a far-field boundary condition is acceptable when $\vec{Q}^{(o)}$ represents freestream values of the flow variables, and is only strictly valid for one-dimensional flow. In most applications, either the far-field boundary will be sufficiently distant from the configuration that $\vec{Q}^{(i)} \simeq \vec{Q}^{(o)}$, or if not, $\vec{Q}^{(i)}$ will be a better approximation of the desired outflow boundary condition than $\vec{Q}^{(o)}$. However, in an OVERFLOW-CHARM simulation, neither of these conditions apply. Consider the outflow boundary just behind the rotor blade in an OVERFLOW-CHARM simulation. When the standard characteristic boundary condition is applied after the CHARM boundary condition, the tangential component of the CHARM-induced velocities will be overwritten by those extrapolated from the internal cell. However, the values of the internal cell depend on the boundary condition applied at the previous step, for which the tangential velocities have also been extrapolated. This results in an outflow bound-

ary which has a very weak correlation to the flow conditions computed by CHARM, which is accounting for the full wake geometry. The velocities tangential to this boundary are significant for the prediction of the rotor wake flow physics, as the tip vortex- and vortex sheet-induced velocities will be predominantly tangential to this boundary.

It is not sufficient to simply remove the application of the characteristic boundary condition for OVERFLOW-CHARM simulations, because non-physical pressure wave reflections from the CFD domain boundary can have a non-negligible impact on the rotor loads. While the standard characteristic boundary condition has a negative impact on the boundary velocities, it does mitigate these pressure wave reflections.

To address these two issues, a new free-wake characteristic boundary condition (F-W CBC) was devised that both preserves the CHARM-induced velocities on the CFD boundary and treats pressure wave characteristics appropriately to mitigate their reflection back into the CFD domain. This is achieved by setting

$$\text{sign}(\mathbf{\Lambda}) \equiv \begin{bmatrix} -1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & -1 \end{bmatrix}. \quad (9)$$

in equation 3. This ensures that all characteristics are treated the identically at the OVERFLOW-CHARM boundary and achieves the desired effect on the vorticity contours and pressure waves. Table 1 breaks the motivation for this new free-wake characteristic boundary condition down into its simplest terms.

Table 1: Quality of characteristic treatments for various implementations of the CHARM boundary condition.

	Correlation b/w Boundary Vel. and Induced Vel.	Ability of BC to Prevent Wave Reflection
No CBC	Good	Poor
CBC	Poor	Good
F-W CBC	Good	Good

4.4. Boundary Value Interpolation

Because the CHARM solution is updated much less frequently than the OVERFLOW solution, CFD boundary values may change significantly when the CHARM wake is updated. This has been observed to cause large, non-physical oscillations in the blade loads at the CHARM update frequency, which become more pronounced as the CFD boundaries approach the blade.

To mitigate this, velocities and pressures can be interpolated between CHARM update steps. For example, if the CHARM solution updates every five degrees, at $\psi = 7.5$ degrees, the boundary values will be the average of the values computed at $\psi = 5$ degrees and those computed at $\psi = 10$ degrees. Since the CHARM model will not have reached ten degrees yet, the values at ten degrees need to be predicted. There are two methods that have been investigated to predict these values.

4.4.1. Velocity-Predictive Interpolation (VPI)

The CHARM-induced boundary values which need to be predicted can be set to those computed during the previous revolution, assuming periodicity of the rotor aerodynamics. This reduces the discontinuity in boundary values before and after CHARM update steps, and thus, the quality of OVERFLOW-CHARM results. While pressure is derived from $\frac{\partial\phi}{\partial t}$, because $\frac{\partial\phi}{\partial t}$ can vary significantly from step to step at a given point, interpolating $\frac{\partial\phi}{\partial t}$ directly can cause non-physical pressures at the boundary, so pressure is interpolated instead. The forces on the blades are computed and sent to CHARM at the rate of CHARM solution updates, as normal. Referring to Figure 2, the flow conditions at an arbitrary azimuth angle $U(\psi)$ are computed as:

$$U(\psi) = aU(\psi_1) + (1-a)U(\psi_2), \quad (10)$$

where a is the interpolation weight of the first term. The red text indicates that $U(\psi_2)$ is being predicted and the blue text indicates that $U(\psi_1)$ is computed instantaneously at CHARM update steps. The prediction for $U(\psi_2)$ is:

$$U(\psi_2) = U(\psi_2 - 360), \quad (11)$$

where the green text indicates that $U(\psi_2 - 360)$ comes from saved data.

4.4.2. Force-Predictive Interpolation (FPI)

When the CHARM-predicted rotor wake is insufficiently periodic due to unconverged blade loads or its inherently unsteady nature, the VPI method still results in discontinuities in boundary values at the CHARM update step. The magnitude of the discontinuity is proportional to the aperiodicity of the CHARM wake. This discontinuity indicates that there are still oscillations in the predicted blade loads and CFD flow field at the CHARM update frequency, though the oscillations are weaker than with no interpolation at all. To address this issue, a new method for predicting velocities was devised. Instead of using boundary values computed during the previous revolution, the CHARM

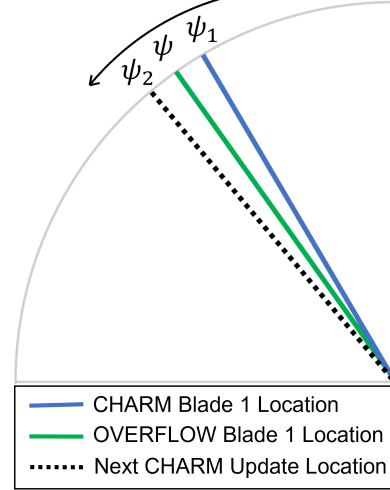


Figure 2: Illustration of definitions for ψ , ψ_1 , and ψ_2 in Equations 10 and 11.

model leads the CFD solution by one update step. Since CHARM needs the blade lift distribution to generate the wake at the next update azimuth angle, the loads are predicted using those computed during the most recent blade passage through that azimuth angle. Boundary values can then be interpolated exactly, since the induced conditions at CHARM update angles before and after the current angle are both known. This results in an interpolation between computed and predicted boundary values that is always smooth. The trade-off is that the CFD solution “lags” the CHARM solution by one blade passage, which is acceptable for periodic cases. Referring to Figure 3, the flow conditions at an arbitrary azimuth angle $U(\psi)$ are computed as:

$$U(\psi) = aU(\psi_1) + (1-a)U(\psi_2). \quad (12)$$

No terms are being predicted in this equation, ensuring that there are no discontinuities at the update steps. Instead, the blade lift distribution at the next update step $F_{blade1}(\psi_2)$ is predicted as:

$$F_{blade1}(\psi_2) = F_{blade2}(\psi_2). \quad (13)$$

The advantages of V-Predictive Interpolation and F-Predictive Interpolation are illustrated in Figure 4.

Boundary value interpolation should not be used when there is aperiodic motion of the CFD grid outer boundaries, such as during convergence of an aeroelastic simulation. There will be a much more significant discontinuity in boundary values when using VPI, since the induced velocities and pressures were computed at a different point in space during the previous revolution. FPI uses the point locations from the previous revolution to compute the forward induced conditions, which will be incorrect in the case of aperiodic grid boundary motion. This can be circumvented by intro-




Figure 3: Illustration of definitions for ψ , ψ_1 , and ψ_2 in Equations 12 and 13.




Figure 4: Interpolation quality without interpolation, with Velocity-Predictive Interpolation (VPI), and with Force-Predictive Interpolation (FPI).

ducing a bounding grid that has prescribed periodic motion onto which CHARM-induced conditions can be applied. This ensures that boundary point locations are identical from revolution to revolution.

4.5. Single Gridded Blade (SGB) Simulation

Originally, OVERFLOW-CHARM required all blades to be resolved with CFD grids. In flight conditions where quick convergence to periodicity is expected, it may be desirable to resolve only a single blade with a CFD grid. This capability was added to OVERFLOW-CHARM so that the loads and deformations of ungridded blades are duplicated from those computed on a single gridded blade based on assumed periodicity of the solution. CHARM still predicts the wake from the true number of blades in this scenario so that correct blade-wake interactions are predicted.

This approach reduces the size of the CFD grid by

a factor of the number. However, computational cost savings may be slightly less pronounced since parallel efficiency begins to drop when the CFD grid becomes very small. Also, blade loads and geometries at each azimuth station in the CHARM simulation are updated only once every revolution instead of once every blade passage, potentially reducing the convergence rate of the OVERFLOW-CHARM solution. In axisymmetric cases such as hover, this latter drawback is circumvented by setting the blade loads and properties of all blades equal to those computed on the resolved blade, making single-blade simulations very cost-effective for hover analysis.

5. Validation Cases

5.1. Case 1: Vortex Advection

This case is intended to validate OVERFLOW-CHARM's ability to accurately model the simple flow physics of vortices advecting through a control volume. The vortices are initialized using the induced flow-field of a CHARM simulation of a single subscale S76 rotor blade in a region just below the outer radius of the rotor plane, as illustrated in Figure 5. This initial flow field is then predicted using OVERFLOW-CHARM, where the boundary conditions of the control volume are those induced by the rotor blade wake simulated by CHARM, and the interior flow is computed by OVERFLOW.

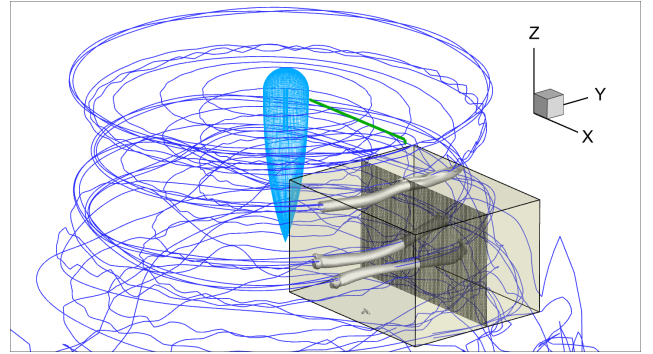


Figure 5: Setup of vortex advection scenario. OVERFLOW region boundaries are colored by Q-criterion, CHARM vortex wake is in blue, and the rotor blade is in green. The shown slice is used for visualization.

The computational domain is a single $100 \times 100 \times 101$ box grid which is static in space below the initial rotor blade position. It extends from $0.54R$ to $1.43R$ in the blade spanwise direction, $-0.41R$ to $0.30R$ in the blade chordwise direction, and $-0.71R$ to $-0.18R$ in the vertical direction. Grid points are clustered around the path of the blade tip vortices. Blade

forces are pre-computed and do not change as the simulation progresses to focus analysis on the induced flow field within the region of interest.

Unless otherwise noted, the numerical options for all simulations are the same. Time advancement is performed using Euler implicit dual time stepping with 39 subiterations and a constant CFL number of 10. The time step corresponds to 0.25 degrees of rotation of the rotor blade, which translates to approximately 3000 steps for a vortex to advect through the computational domain. The right-hand-side terms were discretized using the fourth-order Rho upwind spatial differencing with 4/2 TLNS3D dissipation where the fourth- and second-order dissipation constants were 0.04 and 2.0, respectively. The left-hand-side terms were discretized using an ARC3D diagonalized Beam-Warming scalar pentadiagonal scheme with fourth-order TLNS3D dissipation. The one-equation Spalart-Allmaras turbulence model was used with a rotational correction. The reference condition was at the rotor tip, with a Mach number of 0.65 and a Reynolds number per inch of 120,000. The CHARM vortex wake is resolved with 360 vortex elements per filament and is updated every five degrees, corresponding to five turns of rotor wake resolved with full-span vortex filaments. Root and tip filaments are preserved for an additional five revolutions.

Because vortex pairing is predicted by CHARM, there is strong aperiodicity in the CHARM solution at all times. Applying VPI to this case results in very strong oscillations in the solution at the CHARM update frequency, enough to impact the advection rate of the vortices. When FPI is applied, however, these oscillations nearly disappear. Figure 6 compares the pressure field predicted when using VPI to the pressure field when using FPI. With VPI, the pressure field is utterly dominated by the oscillations resulting from the jump in boundary values at the update frequency. With FPI, on the other hand, the true dynamics of the flow field are clearly visible from the pressure field.

Three different characteristic boundary conditions are applied to the CFD boundary: no Characteristic Boundary Condition (CBC), a standard CBC, and the Free-Wake CBC. Figure 7 compares the pressure fields when employing the Free-Wake CBC and when not applying any CBC. Most notable are the nonphysical pressure gradients along the right-side boundaries predicted when not applying a CBC, caused by the discrepancy between the internal and boundary pressures predicted by OVERFLOW and CHARM, respectively, which is not resolved by the boundary condition.

While the pressure discrepancy can be resolved with a standard CBC, it is necessary to inspect the velocity field to understand why the Free-Wake CBC is advan-

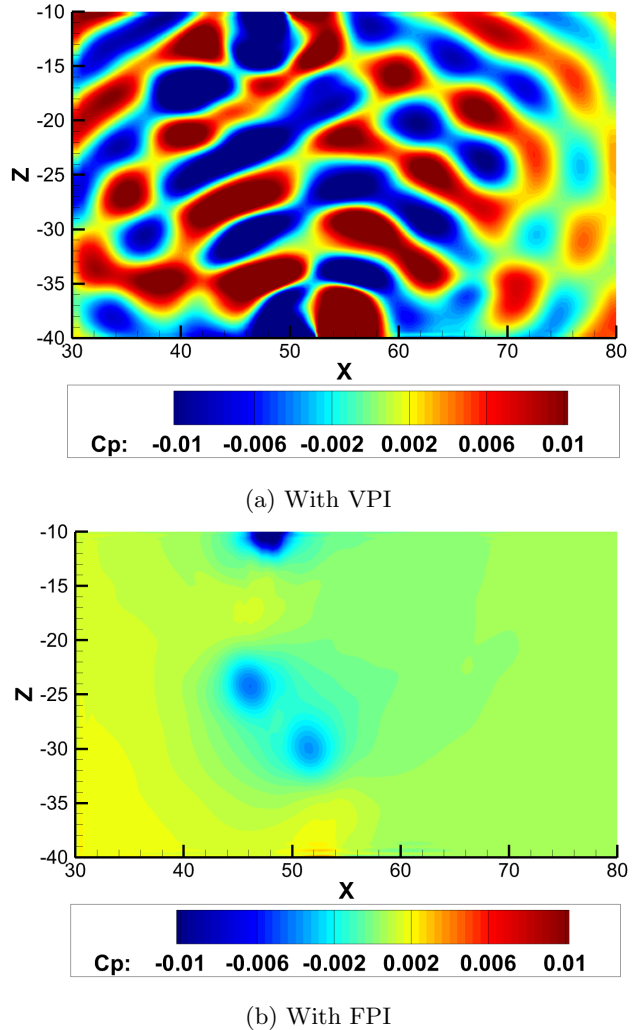


Figure 6: Comparison of pressure fields on a $y = 0$ slice for the vortex advection scenario when using Velocity-Predictive Interpolation (VPI) vs. Force-Predictive Interpolation (FPI) at $\psi = 270^\circ$.

tageous. In Figure 8, isocontours of radial momentum are compared between a standard CBC (red) and the Free-Wake CBC (black). In the first image, there is a region of outflow on the bottom boundary that triggers velocity extrapolation by the standard CBC. This causes a discrepancy in the CHARM-predicted velocity and the actual velocity which is applied on the boundary. In the second image, where the outflow region is no longer present, the boundary radial velocity at that location changes discontinuously to the CHARM-induced value instead of the extrapolated value. The standard CBC formulation creates a boundary condition where the velocity tangential to the boundary depends on the sign of the velocity normal to the boundary, which is not physically consistent with the fact that the velocity is being predicted by CHARM. This issue is resolved by the Free-Wake CBC, where the behavior of the black radial velocity contours in Figures

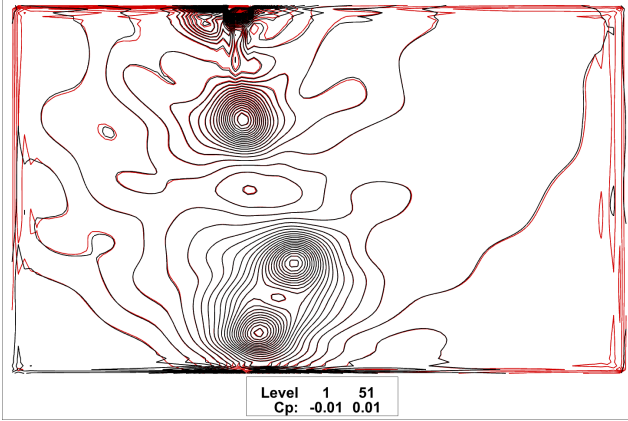
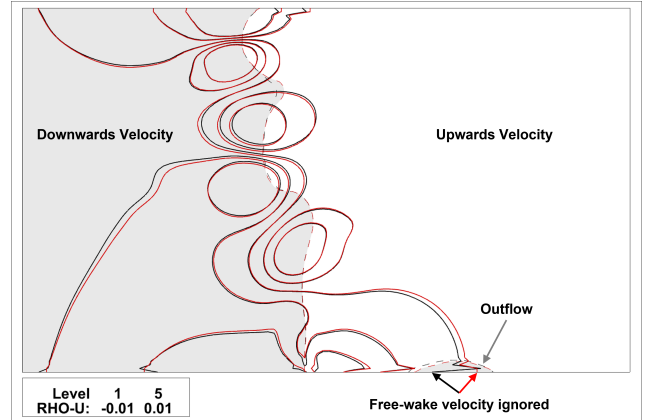


Figure 7: Pressure isocontours for the vortex advection scenario with and without a Free-Wake CBC applied (black and red, respectively) at $\psi = 177.5^\circ$.

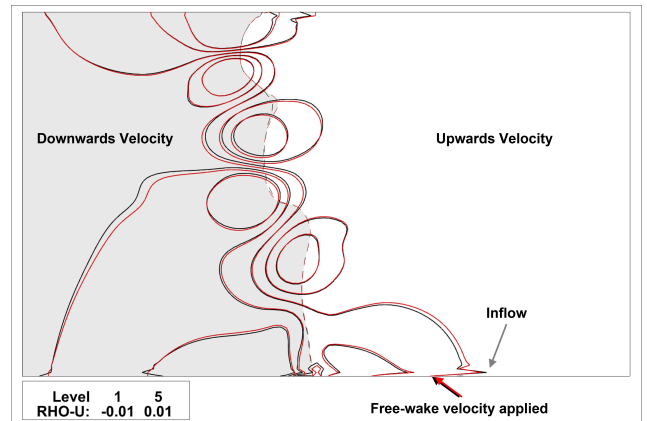
8a and 8b, does not depend on the sign of the normal velocity at the boundary.

The oscillations in the velocity near the boundaries, observed regardless of the type of CBC applied, depend on the spatial differencing scheme. When computing gradients with central-differences, these oscillations propagate much farther into the domain than for upwind schemes. First- or second-order upwind schemes exhibit far less pronounced oscillations than third- or fourth-order upwind schemes, as illustrated in Figure 9. However, the application of a second-order scheme on the same computational grid results in much greater numerical dissipation of vorticity, resulting in much weaker vortices even after a short time, as observed in Figure 10 through the nondimensional Q-criterion contour magnitudes.

In light of these observations, it is recommended to use high-order upwind schemes for hybrid CFD simulations to produce accurate results on a cost-effective grid. The sharp gradients near the boundary are likely present because the boundary condition is not time-accurate, in the sense that interpolating the free-wake-induced velocities between CHARM update steps is not physical. However, when employing upwind schemes, these oscillations do not propagate into the flow field. The same result is expected when employing central-difference schemes on a non-uniform computational grid. Of course, such behavior would likely be present for any boundary condition on a computational grid this small, and based on the results of previous studies affirming the efficacy of hybrid CFD techniques [9, 11, 13, 14], these phenomena do not detract from the value of the overall methodology. This is especially true considering the extremely low cost of the simulations in this study. This cost of each simulation in this study was approximately 115 CPU-hours, costing six seconds per iteration over 1440 steps. Simulations were run on



(a) $\psi = 282.5^\circ$



(b) $\psi = 307.5^\circ$

Figure 8: Isocontours of radial (rightward) momentum for the vortex advection scenario with a Free-Wake CBC (black) and a standard CBC (red) applied. The white and grey regions delineate positive and negative vertical velocity, respectively (slight variation between the two cases, as indicated by the dashed lines).

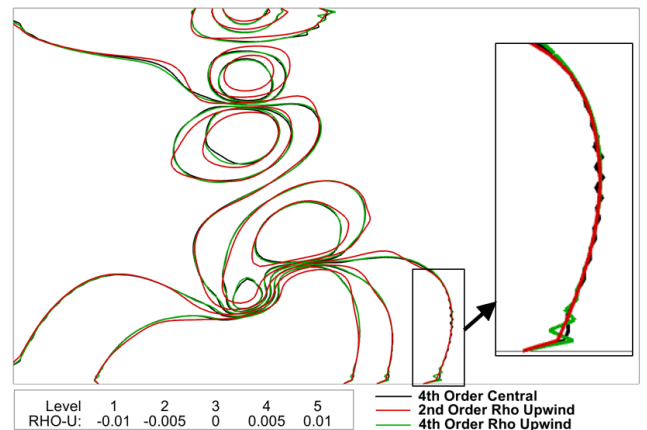


Figure 9: Isocontours of radial (rightward) momentum for the vortex advection scenario with various spatial differencing schemes applied at $\psi = 12.5^\circ$.

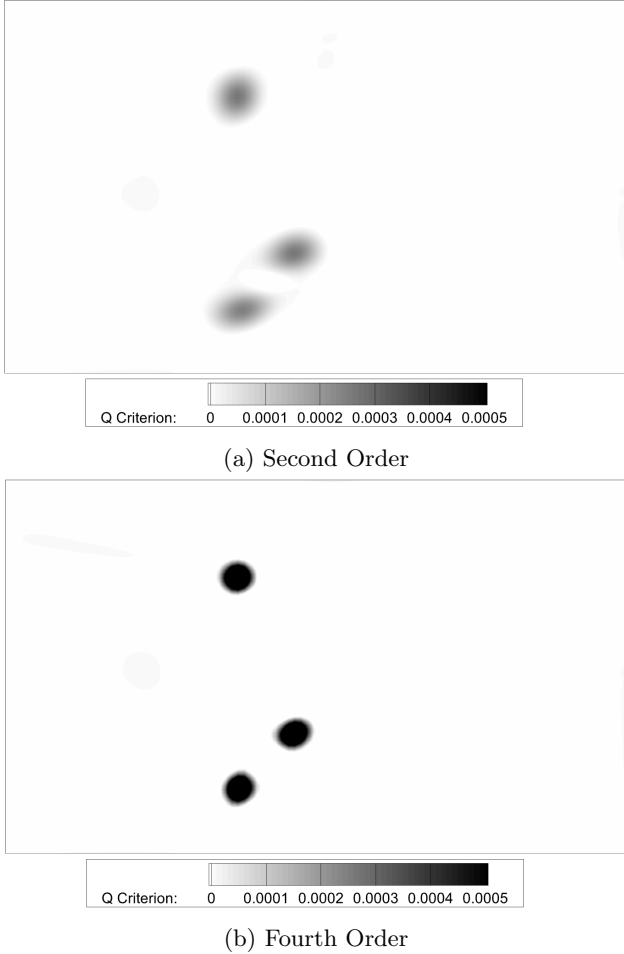


Figure 10: Contours of nondimensional Q-criterion for the vortex advection scenario with different orders of accuracy of the Rho upwind scheme at $\psi = 12.5^\circ$.

48 MPI processes on the US DOD HPCMP machine “Gaffney” with Intel®Xeon®Platinum 8168 CPUs at 2.7 GHz. The low cost is partially because only the region of interest need be resolved with a CFD grid and also that the rotor wake is fully initialized before the simulation begins. To produce similar results using a conventional CFD simulation would require the rotor blades and full flow field be resolved with a dense CFD grid and multiple revolutions of expensive computation to initialize the rotor wake. This would almost certainly be orders of magnitude more expensive than the OVERFLOW-CHARM simulations presented here.

5.2. Case 2: Wing-Integrated Propeller

The second case, a subscale propeller mounted on a wing, is well-suited to demonstrate the improvements in rotor/blade tracking and the influence of the unsteady free-wake pressure. The simulation analysis is illustrated in Figure 11, where a small region around the propeller is resolved with an OVERFLOW grid, and propeller wake, nacelle, and wing aerodynamics are resolved in the CHARM model. A more detailed description of this configuration and more extensive simulation results can be found in the literature [13]. The results presented here focus on the zero angle of attack case at a freestream Mach number of 0.083 and a propeller tip Mach number of 0.375.

Time is advanced with dual time-stepping and twenty second-order Newton subiterations at a time step corresponding to a quarter-degree rotation of the propeller. Fourth-order central spatial differencing and an ARC3D Beam-Warming block tridiagonal scheme are applied for the right- and left-hand side terms, respectively. The $k\omega$ -SST RANS turbulence model was employed. The CHARM wake model represents the rotor wake with two revolutions of full-span vortex elements and five revolutions of root and tip vortex elements, updated 72 times per revolution. The nacelle and wing are resolved using doublet panels in CHARM, and at the sharp trailing edge of the wing, panels are shed in the CHARM model to capture the lift generated by the wing using a Morino panel method [15]. The wing lifting panels ensure that the lift generated by the wing is accurately reflected in the CHARM model and correctly influences the OVERFLOW domain boundary conditions.

In this case, the rotor axis is no longer parallel to the z-axis of the inertial frame, and the blades are rigid, thin, and highly twisted. This configuration would have been very unintuitive to model using the original OVERFLOW-CHARM formulation, whereas with the improvements to rotor and blade geometry tracking, very little additional effort is required to configure




Figure 11: Wing-integrated propulsion case setup. CFD region is in grey, CFD surfaces are in green, and CHARM-resolved surfaces are in blue.

the OVERFLOW-CHARM input files compared to a conventional OVERFLOW simulation. The simulation can also be performed in the inertial frame for CFD of a fixed-wing system (with x in the freestream direction), which was not possible before.

The accuracy of aerodynamic predictions is also positively impacted by the improvements discussed in the previous section, particularly the unsteady free-wake pressure applied to the CFD boundary. To establish this, propeller thrust predictions employing isentropic boundary pressure and unsteady free-wake pressure are compared to the thrust predicted an OVERFLOW simulation with a conventionally-sized grid (25.7 million points, see Reference [13] for more details). When the boundary pressure and density are simply approximated using isentropic relations, OVERFLOW-CHARM predicts the mean propeller thrust coefficient to be 2% lower than that predicted by the conventional OVERFLOW simulation. However, when the pressure and density are instead derived from the unsteady free-wake pressure predicted by CHARM, OVERFLOW-CHARM predicts the mean propeller coefficient within 0.3% of the OVERFLOW prediction. See Table 2 for the thrust predictions of each method. Because the cost of CHARM-related routines comprise less than 1% of the computational cost of an OVERFLOW-CHARM simulation, the additional cost of computing the unsteady potential function is negligible when considering the increase in accuracy that it provides.

Because the OVERFLOW-CHARM solution for this case converges to periodicity, the difference in the converged blade loads between VPI and FPI is not significant, however, FPI improved the quality of the solution during convergence of the simulation, as seen in Figure 12. This may be advantageous if a very short OVERFLOW-CHARM simulation is desired for quick turn-around applications where strong convergence is not required. The remaining oscillations ob-

Table 2: Propeller thrust coefficient as predicted by OVERFLOW, OVERFLOW-CHARM with isentropic boundary pressure, and OVERFLOW-CHARM with unsteady free-wake boundary pressure

	Thrust Coeff.
OVERFLOW	0.23048
OVERFLOW-CHARM (Isentropic Pressure)	0.22644 (-1.91%)
OVERFLOW-CHARM (Unsteady F-W Pressure)	0.23151 (+0.29%)

served in the FPI solution are likely the result of slight differences in pressure predictions between CHARM and OVERFLOW at the boundary, which is unavoidable due to the different physical formulations of the two methods. It is possible that using OVERFLOW-computed velocities to convect CHARM vortex elements within the CFD domain may remedy this effect, though validation of this capability still ongoing.

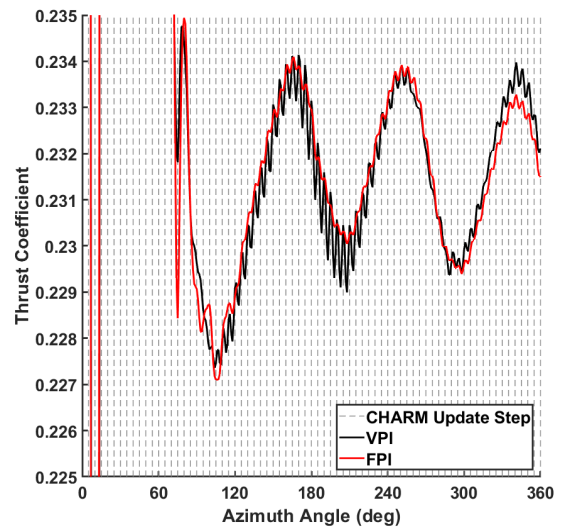


Figure 12: Comparison of propeller thrust coefficient variation at the beginning of an OVERFLOW-CHARM simulation with Velocity-Predictive Interpolation (VPI) and Force-Predictive Interpolation (FPI). CHARM update steps are shown for reference.

5.3. Case 3: Rotor in Hover

A rotor in hover is the ideal scenario to demonstrate the Single Gridded Blade (SGB) methodology as not only is it periodic, it is axisymmetric. Therefore the non-gridded blades in the CHARM model can use the instantaneous blade loads from the gridded blade instead of those previously computed at a specific azimuth angle. When SGB simulations are run in axisymmetric

mode, the convergence rate should be just as fast as if all blades were gridded. Therefore there is no trade-off for the increased cost savings SGB simulation provides in this case. The simulation setup for a rotor in hover using the SGB method is depicted in Figure 13.

The rotor properties are given in Table 3. The blades were untwisted and untapered at a constant collective angle of 12 degrees. The rotation speed of the rotor was 50 Hz, corresponding to a tip Mach number of 0.08.

Table 3: Table of rotor properties for the micro-scale rotor.

Rotor Radius	Blade Chord	Root Cutout
91.6 mm	21 mm	16 mm
Blade Airfoil	Thickness	Camber
Circular Arc	3.7%	3.3%

Time-advancement was performed with dual time stepping using up to sixty second-order subiterations, a local time step factor of 0.1, and a minimum CFL number of five. Subiterations were stopped when the residual dropped two orders of magnitude. The physical time step was selected so that the rotor advances 0.25 degrees per time step. The right-hand-side terms were discretized using the fourth-order Rho upwind spatial differencing with 4/2 TLNS3D dissipation where the fourth- and second-order dissipation constants were 0.04 and 10.0, respectively. The left-hand-side terms were discretized using an ARC3D diagonalized Beam-Warming scalar pentadiagonal scheme with fourth-order TLNS3D dissipation. Viscous terms were included in all directions given the relatively low Reynolds number. The one-equation thin-layer Spalart-Allmaras turbulence model provided turbulence closure.

Employing the SGB method, predicted thrust and power is within 0.4% of a non-SGB simulation at 56% the computational cost. This is consistent with the theoretical performance of the SGB method in axisymmetric cases, where accuracy should not be significantly effected and computational cost should be reduced by a factor of the number of blades.




Figure 13: Single Gridded Blade (SGB) simulation setup for a rotor in hover.

To focus the analysis on the Free-Wake CBC, a fully-gridded rotor simulation was performed with a bounding box grid around the rotor blades, as illustrated in Figure 14. This aids in inspecting inflow and outflow regions. Analysis is focused on the region of interest shown in pink, which contains regions of both inflow and outflow on the same boundary, as predicted by CHARM. Figure 15a plots the boundary vorticity predicted by CHARM when no CBC is applied. There are smooth regions of vorticity corresponding to the blade tip vortices and the shed vortex sheet. Figure 15 compares the y-vorticity on the boundary region of interest specified in Figure 14. This illustrates gradients in the velocities tangent to the boundary. The y-vorticity field produced when no CBC is applied appears in Figure 15a, and serves to illustrate what the y-vorticity should look like. The tip vortices from each blade are present, as well as the vortex sheet shed from the blade just in front of the boundary region of interest. When the standard CBC is applied, as observed in Figure 15b, spurious vorticity is generated along the boundary between the inflow and outflow regions due to the difference in treatment of tangential velocities within each region. The Free-Wake CBC, however, correctly reproduces the CHARM-induced vorticity on the boundary, as seen in Figure 15c. It is also the case that the inflow/outflow boundary is modified by the free-wake CBC, indicating that outgoing characteristics are being appropriately treated without compromising the tangential velocity field. In Figure 16, the Free-Wake CBC also has the same mitigating effect on rotor integrated thrust and power oscillations as the standard CBC, indicating that pressure wave reflections from the boundary are being effectively prevented.

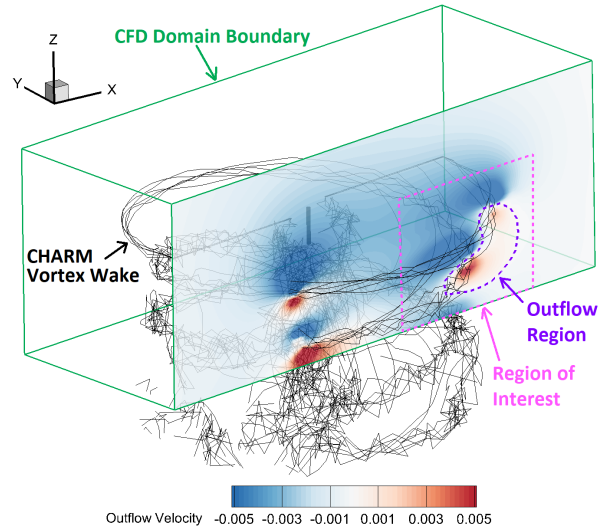
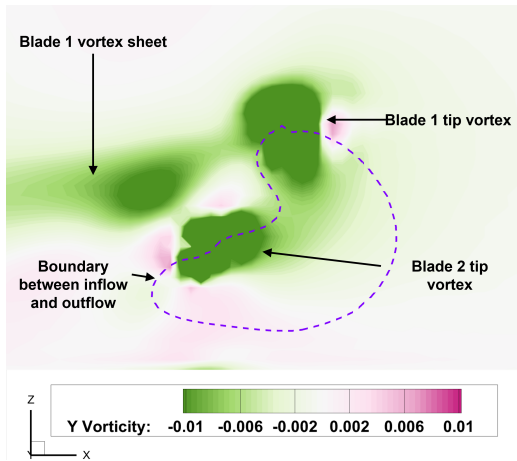
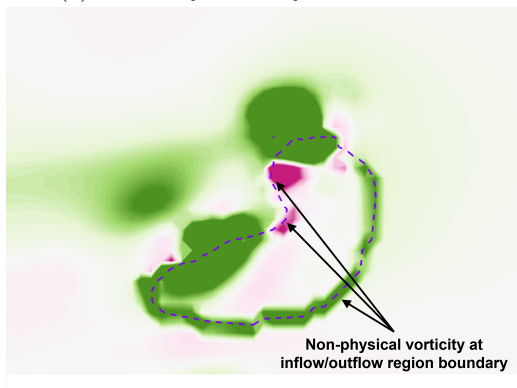


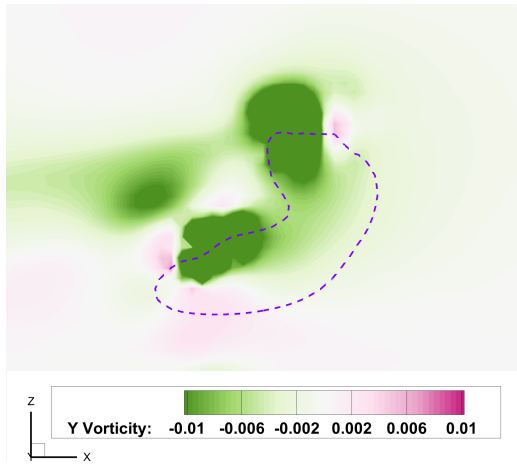
Figure 14: Simulation setup of the CBC demonstration for the rotor in hover with contours of CHARM-induced outflow velocity on the boundary of interest.



(a) Boundary Vorticity With No CBC

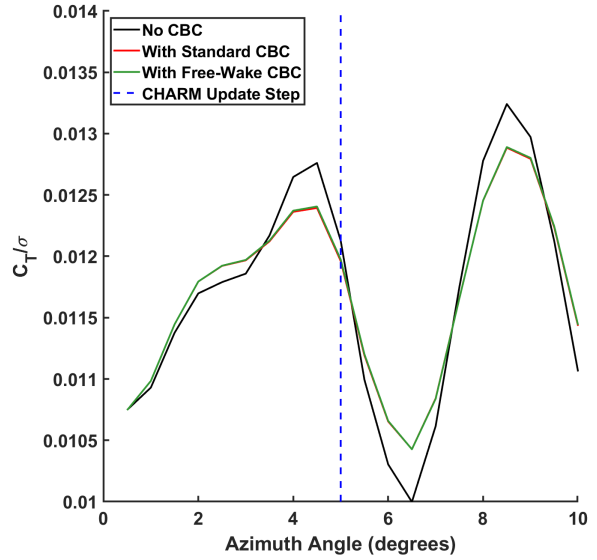


(b) Boundary Vorticity After Standard CBC

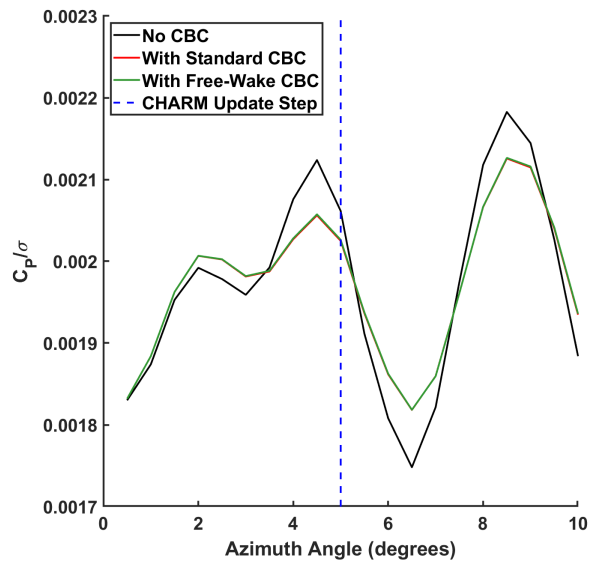


(c) Boundary Vorticity After Free-Wake CBC

Figure 15: Comparison of boundary-normal vorticity contours on the CFD boundary in an OVERFLOW-CHARM simulation before the application of any CBC and after the application of the original and improved Riemann boundary conditions. See Figure 14 for context.



(a) Thrust Coefficient



(b) Power Coefficient

Figure 16: Comparison of boundary-normal vorticity contours on the CFD boundary in an OVERFLOW-CHARM simulation before and after the application of the improved Riemann boundary condition. See 14 for context.

6. Concluding Remarks

The improvements made to the OVERFLOW-CHARM hybrid CFD framework make it uniquely accurate and flexible among other hybrid CFD methods for rotorcraft aeromechanical analysis. In particular, the boundary condition, which serves as the primary interface between the two computational domains, is treated carefully and thoughtfully for maximum accuracy. The inclusion of the unsteady potential function and new interpolation and boundary conditions measurably improve the accuracy of the results in the selected example scenarios. Improvements to the handling of advanced rotor configurations, blade geometry calculation, and SGB capability also make OVERFLOW-CHARM efficient for application to modern rotorcraft analysis. It is recommended that any hybrid CFD coupling between Navier-Stokes and vortex-based solvers follow the best practices established here in their implementation.

7. Future Work

There exists the capability within OVERFLOW-CHARM to apply the OVERFLOW-predicted velocity field to determine the velocity of the free-wake elements in the CHARM model within the OVERFLOW domain. However, it is not possible to make use of this capability when using FPI due to the asynchronicity of the OVERFLOW and CHARM solutions. A alternative solution to this problem should be investigated, perhaps applying a periodic assumption for the OVERFLOW velocity field. This capability would improve OVERFLOW-CHARM's ability to model rotor aerodynamics when separation, strong compressibility, or vortex burst are present within the OVERFLOW domain.

8. Acknowledgements

Funding for this work is provided by the Department of Defense (DoD) Science, Mathematics, and Research for Transformation (SMART) scholarship program. Computational resources are allocated through the DoD High Performance Computing Modernization Program (HPCMP) through S/AAA Meghan Goldsborough.

References

- [1] Anusonti-Inthra, P., "Developments and Validations of Fully Coupled CFD and Particle Vortex Transport Method for High-Fidelity Wake Modeling in Fixed and Rotary Wing Applications," Tech. Rep. NASA/CR-2010-216696, NASA, May 2010.
- [2] Chen, P.-W. and Sankar, L., "A Hybrid Navier Stokes-Free Wake Method for Modeling Tandem Rotors," *Proceedings of the 7th Asian/Australian Rotorcraft Forum*, Nov. 2018.
- [3] Battey, L. S., *A Hybrid Navier-Stokes/Vortex Particle Wake Methodology for Modeling Helicopter Rotors in Forward Flight and Maneuvers*, Master's thesis, Georgia Institute of Technology, May 2018.
- [4] Rajmohan, N., *Application of Hybrid Methodology to Rotors in Steady and Maneuvering Flight*, Ph.D. thesis, Georgia Institute of Technology, Aug. 2010.
- [5] Thomas, S., *A GPU-Accelerated, Hybrid FVM-RANS Methodology for Modeling Rotorcraft Brownout*, Ph.D. thesis, University of Maryland, College Park, Jan. 2013.
- [6] Nichols, R. and Buning, P., *OVERFLOW User's Manual, Version 2.2*, NASA Langley Research Center, Hampton, VA, Aug. 2010.
- [7] Bliss, D. B., Tenske, M. E., and Quackenbush, T. R., "A New Methodology for Free Wake Analysis Using Curved Vortex Elements," Tech. Rep. NASA/CR-3958, NASA, Dec. 1987.
- [8] Quackenbush, T. R., Bliss, D. B., Wachspress, D. A., Boschitsch, A. H., and Chua, K., "Computation of Rotor Aerodynamic Loads in Forward Flight Using a Full-Span Free Wake Analysis," Tech. Rep. 19930014869, NASA, Oct. 1990.
- [9] Jacobson, K., Grubb, A., and Smith, M., "Performance and Physics of a S-76 Rotor in Hover With Non-Contiguous Hybrid Methodologies," *54th AIAA Aerospace Sciences Meeting*, Jan. 2016.
- [10] Wilbur, I. C., Moushegian, A., Smith, M. J., and Whitehouse, G. R., "Complex Vehicle Design and Analysis with Hybrid Methodologies," *Annual Forum Proceedings - AHS International*, Jan. 2018.
- [11] Wilbur, I. C., Moushegian, A., Smith, M. J., and Whitehouse, G. R., "UH-60A Rotor Analysis with an Accurate Dual-Formulation Hybrid Aeroelastic Methodology," *Journal of Aircraft*, Vol. 57, No. 1, Jan. 2020, pp. 113–127.

- [12] Wachspress, D. A., Quackenbush, T. R., and Boschitsch, A. H., “Rotorcraft Interactional Aerodynamics with Fast Vortex/Fast Panel Methods,” *Journal of the American Helicopter Society*, Vol. 48, No. 4, 2003, pp. 223–235.
- [13] Moushegian, A., Smith, M. J., Whitehouse, G. R., and Wachspress, D. A., “Implementation of a Dual-Solver Hybrid Approach for Rotating System Simulations in HPCMP CREATE(TM)-AV HELIOS,” *AIAA Scitech 2021 Forum*, Jan. 2021.
- [14] Moushegian, A., Smith, M. J., Whitehouse, G. R., and Wachspress, D. A., “Hover Performance in Ground Effect Prediction Using a Dual Solver Computational Methodology,” *Annual Forum Proceedings - AHS International*, May 2021.
- [15] Morino, L. and Kuo, C.-C., “Subsonic Potential Aerodynamics for Complex Configurations: A General Theory,” *AIAA Journal*, Vol. 12, No. 2, Feb. 1974, pp. 191–197.